AUSTRALIAN
communications

JUNE 1994 The Networking and Telecommunications Management Magazine $5.50

A SWITCH IN TIME?

Can switched Ethernet plug the gap until ATM arrives?

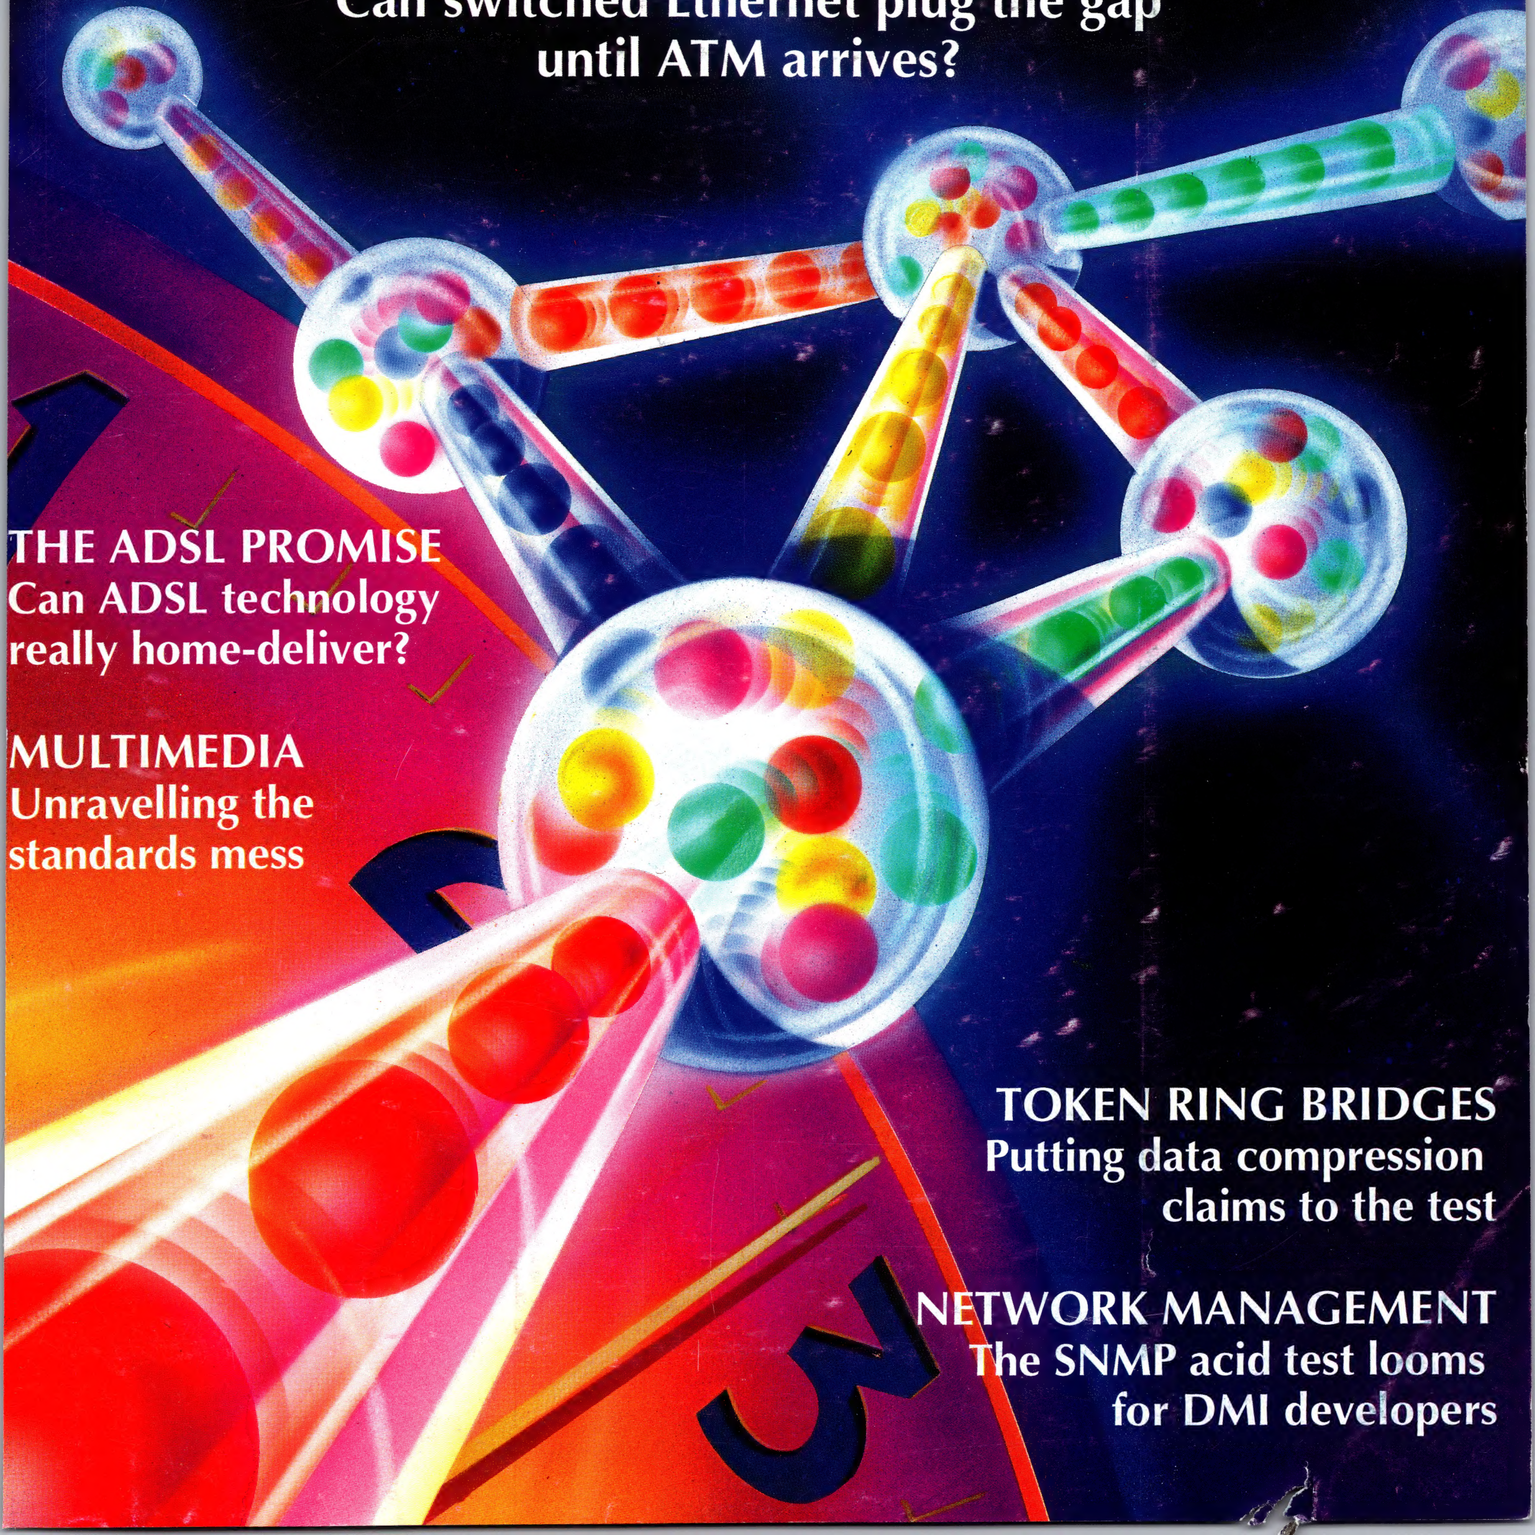

THE ADSL PROMISE
Can ADSL technology really home-deliver?

MULTIMEDIA
Unravelling the standards mess

TOKEN RING BRIDGES
Putting data compression claims to the test

NETWORK MANAGEMENT
The SNMP acid test looms for DMI developers

THE COMPETIT

INTRODUCING THE MMAC-PLUS,™ THE

ION IS FINISHED

FIRST TRUE PACKET SWITCHING HUB.

With the MMAC-Plus, you'll never have to compete for bandwidth again. That's because the MMAC-Plus features integrated switching technology to support both packet-based and ATM cell-based networks. Now you can create dedicated user-per-port LANs that will provide for many classes of service — including E-mail, video conferencing and more — and still guarantee that each user receives the fast, reliable service he or she needs.

Through a unique packet/cell switching backplane and Cabletron's innovative SecureFast™ Packet Switching, the MMAC-Plus delivers the best benefits of ATM technology . . . while preserving your existing investment in Ethernet, Token Ring and FDDI. No other vendor can make that claim.

Features include:

- An aggregate bandwidth in excess of 10 Gbps.
- Aggregate switching greater than 5 million packets/cells per second.
- Over 500 ports of connectivity.
- Support for over 160 switched LAN segments and a virtually unlimited number of shared-access LANs.
- Supports mixed Ethernet, Token Ring, FDDI, ATM and WAN connections.
- Supports integrated MAC layer bridging, network layer routing, SecureFast Packet Switching and ATM cell switching.
- Embedded network analysis tools.
- Fully fault tolerant design.

A hub is not just a hub. For more information on the next-generation hub platform that starts where the competition finishes, call 61-2-950-5900 for your *free*, informative white paper, *Cabletron's MMAC-Plus: A True Switching Hub.*

CABLETRON
SYSTEMS
The Complete Networking Solution™

AUSTRALIA • UNIT 8, ALLAMBIE GROVE BUSINESS ESTSATE • 25 FRENCHS FOREST ROAD • FRENCHS FOREST NSW 2086

AUSTRALIAN communications

Cover: Tony Pyrzakowski

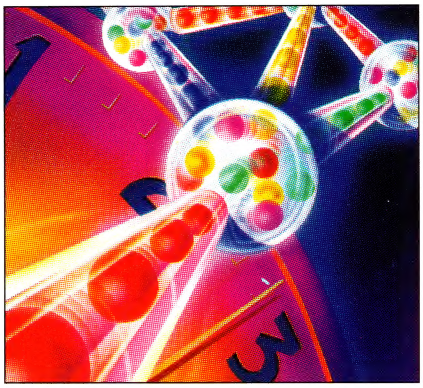

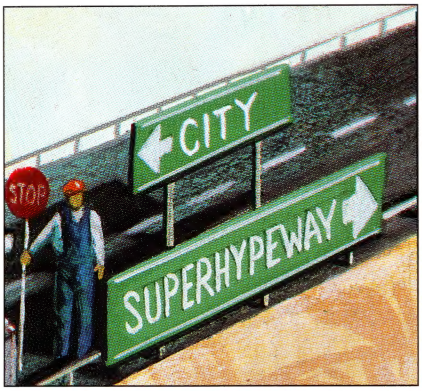

ETHERNET SWITCHING 61

After years of reliable service, Ethernet is stressed out. Powerful network servers, bandwidth-greedy graphics and imaging applications, and skyrocketing numbers of users on each segment are severely taxing many LANs. But now there's an alternative to the technologies typically touted as the answer to overworked Ethernets — FDDI, ATM and fast Ethernet. Recently released Ethernet switches work with conventional cabling and adaptors and deliver 10Mbps to each port. And the potential exists to multiply bandwidth almost indefinitely. In conjunction with the results of hands-on tests of seven of these products, Graeme Le Roux explains the various types of Ethernet switching architecture and how each is employed.

MULTIMEDIA 79

Data superhighways won't carry a whole lot of traffic until a new generation of multimedia applications appears. But what exactly is multimedia and what are the roadblocks slowing down its development? Such is the variety of development it's difficult at this early stage to arrive at conclusions, but one obvious emerging trend is that multimedia has ceased to be a local concept — it has now absorbed the concept of 'convergence' and has metamorphosed by incorporating global networking and broadcasting. It is also increasingly associated with interactivity, and it's here that telephony and broadcasting come into their own. Stewart Fist marches through the multimedia maze to examine early standardisation efforts.

ADSL 87

In our cover story last month, Robin Whittle delivered one of the most comprehensive descriptions of Asymmetrical Digital Subscriber Line (ADSL) technology yet published. This month, in the second installment of his article, Robin builds on that description by delving more deeply into ADSL's inner workings and the enormous technical challenges it must fully overcome. In the labs, the technology promises much, but it's still early days and there are many issues — particularly the problem of cost — to be resolved before ADSL is widely implemented in telco networks around the world. Meanwhile, trials by Telecom and other carriers are being watched with great interest.

JUNE 1994

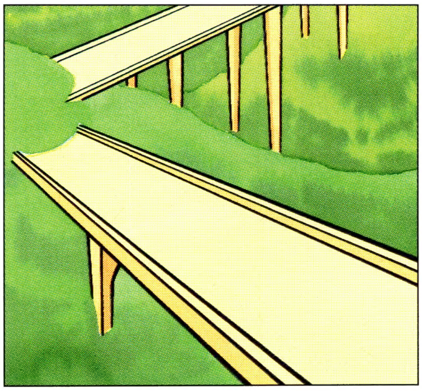

TOKEN RING BRIDGES 97

What are the two most expensive features of any remote Token Ring bridge? The RISC processors? The redundant ports? The flashy LED displays? Guess again. The high cost of WAN bandwidth and the often expensive administrative charges associated with remote connectivity are the real big ticket items. Because of this, it's no big surprise that network managers want remote bridges that can make the most of links between sites. A hands-on test of compression performed by remote Token Ring bridges shows that they can put the squeeze on leased line costs. But beware: vendor compression claims are not always what they seem.

ANALYSIS

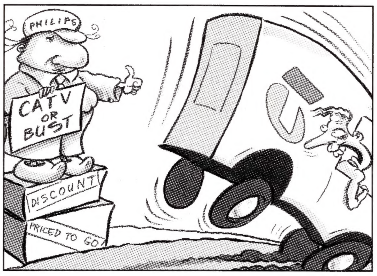

19 Getting a Pay TV Toehold
Philips has won a key role in Telecom's ambitious plans to wire-up most of Australia to receive cable television — but at what cost?

20 Telecom Discovers Coax
21 Stock Watch
22 Telstra's Overseas Drive
24 Optus Opens Up
25 The Packer Purchase
26 Telecom Falls Out of its COT
27 Life After Davey

TECHNOLOGY

31 DMI's SNMP Challenge
Desktop Management Interface developers are judging the merits of two schemes that will let DMI share data with SNMP consoles.

32 Market Watch
34 How Secure is Your LAN?
36 Full Steam Ahead For Full-Duplex
37 Technology Update
42 LAN Views — Graeme Le Roux
43 Open Systems — Alan Lloyd

PRODUCT LEADERS

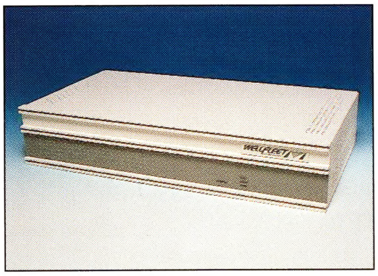

47 Travelling With Magellan
Nortel's new switch combines cell and frame-based trunking to support voice and data networking.

48 A Flexible Remote Router
Wellfleet's Access Node can provide users at branch offices with a range of routing options.

49 Madge-ic Token Ring
Madge's Smart 16/4 AT Plus Ringnode Token Ring adaptor is a high speed option for NetWare users.

OPINION

58 Tom Amos reflects on the social aspects of telecommunications. Is it true that one million Australians can't afford their own telephone?

59 Trish Benson, Co-ordinator of the Consumers' Telecommunications Network discusses the importance of universal service provision.

INTERVIEW

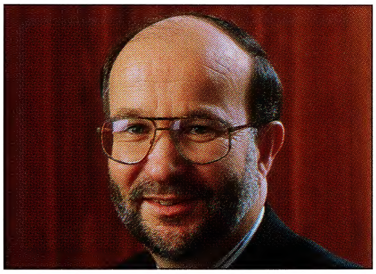

51 Gerry Moriarty
Formerly Assistant Managing Director of the Australian Broadcasting Corporation, Gerry Moriarty is now Managing Director, Network Products, for Telstra and closely involved with many high-profile projects.

LEGAL LINE

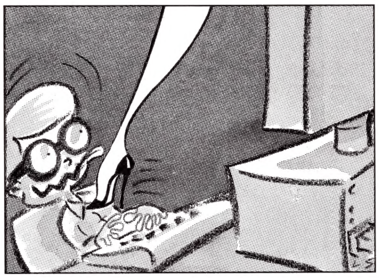

55 The Dial-Up Dilemma
Dial-up bulletin boards are now widely accessible to the public and carry everything from the latest freeware to pornography. Melinda Wallman discusses their regulation.

DEPARTMENTS

VIEWPOINT	6
LETTERS	8
UPDATE	12
PRODUCT FILE	107
THE ATUG PAGES	117
CALENDAR	122

Switching to Ethernet

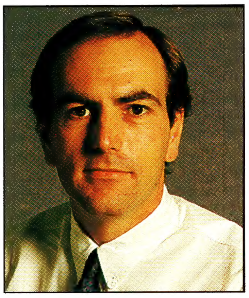

Despite the amount of attention rightly paid to new networking technologies like ATM it's salutary to note that many vendors are planning their new generation network products around a technology now over 20 years old. That technology, Ethernet, has seen off the challenges of pretenders like Token Ring and FDDI and continues to provide the basis for the majority of data networks in the world today. As well as its robustness, the key to Ethernet's longevity has been its relative simplicity and hence its adaptability to new media and new network topologies, as highlighted in this month's edition.

Starting on page 61 (see 'Making the Switch to Ethernet'), Graeme Le Roux provides a run-down on the various Ethernet switching technologies now being rolled out and the important differences between them. In conjunction with Graeme's article, we also publish this month a hands-on look at several leading Ethernet switching products (see 'Ethernet Switches Evaluated' starting on page 63). As Graeme points out, these products have the capability to provide network managers with some much-needed breathing space while we all wait for ATM to truly arrive. This is good news for most vendors, but it may be a little unpalatable for those investing heavily in new fast and full-duplex Ethernet technologies (see 'More Niche Than Necessity?' starting on page 36). While these newer technologies will undoubtedly have a place in future networks, it may not be as soon as some would have us believe.

M Smeaton

In the July edition of *Australian Communications* . . .

MOBILE DATA

With the recent launch of BellSouth Mobile Data and Telecom Mobile Data, Australia now has two public mobile data networks. In next month's edition we take an in-depth look at the technology, tariffs and service offerings of both.

POINT-TO-POINT PROTOCOL

No matter how many LAN protocols two devices have in common, they simply must use the same WAN protocol to work together over the wide area. In our next edition we examine PPP, the increasingly preferred way to tie nets together.

UNIX-TO-SNA GATEWAYS

The backbone protocol debate rages on: SNA champions say APPN developments will ensure its dominance but TCP/IP adherents remain committed to Unix. Coexistence is the only outcome, so next month we examine Unix-to-SNA gateways.

GLOBAL MOBILE COMMUNICATIONS

When two years ago we first reported at length on the various global mobile communications schemes then proposed, many observers doubted they could ever get off the ground. Next month we provide an update on the new space race.

Advertising Enquiries:
Craig Burkill or Dawn Switzer
(02) 264 2200

AUSTRALIAN
communications

EDITOR
Mark Smeaton

CONTRIBUTORS
Tom Amos, Stewart Fist, Liz Fell,
Peter Leonard, Peter Waters,
Stuart Corner, Graeme Le Roux,
Bernard Levy,
Alan Lloyd, Peter Purton

ASSISTANT EDITOR
Sarah Parkes

PRODUCTION MANAGER
Elisabeth Alexander

ILLUSTRATIONS
Raphaella La Roche, Aldo Venturi
Loui Silvestro

SALES AND MARKETING DIRECTOR
Craig Burkill

DEPUTY ADVERTISING MANAGER
Dawn Switzer

CHIEF FINANCIAL OFFICER
Jennifer Stibbard

Subscriptions
Australian Communications is available by paid subscription for $54 per annum. Subscribers receive eleven issues per annum.

Overseas Rates by Air	1 Year	2 Years
New Zealand, Papua New Guinea	$66	$110
Singapore, Indonesia, Malaysia, Brunei, Pacific Islands	$72	$122
Asia — including Hong Kong, India, Korea, Japan, Taiwan	$78	$133
Europe, North America, Middle East, South America	$92	$168

Publishers
Published and distributed monthly by Ostasun Pty. Ltd. (ACN 003 606 102) under licence from ACP Computer Publications, a division of ACP Publishing Pty. Ltd. (ACN 053 273 546) of 54 Park Street, Sydney NSW 2000. ISSN 0818-9021. **Address:** Level 4, 541 Kent Street, Sydney NSW 2000. **Tel:** (02) 264 2200 **Fax:** (02) 264 2244.

Printed at **Offset Alpine Printing**, Derby & Wetherill Streets, Silverwater NSW 2141.

All material contained within *Australian Communications* is protected under the *Commonwealth Copyright Act 1968.* No material may be reproduced in part or whole without prior written consent from the copyright holders.
The opinions expressed in this magazine are not necessarily those of the publisher. We welcome feature articles from contributors. Manuscripts should be typewritten and all text, photographs and illustrations must be accompanied by a stamped, self-addressed envelope stamped to the appropriate value. Retain duplicates of both text and illustrative material; Ostasun (Pty. Ltd.) does not accept responsibility for damage to, or loss of material submitted for publication.
Articles reprinted in this issue from *Data Communications* (USA) are copyright 1994 by McGraw-Hill Inc. All rights reserved. Published from *Data Communications* with the permission of McGraw-Hill Inc., 1221 Avenue of the Americas, New York, New York. Reproduction in any manner, in any language, in whole or in part without the prior written permission of McGraw Hill Inc. is expressly prohibited. Articles reprinted in this issue from *Telecomeuropa* are copyright 1994 by News Bureau Publications. All rights reserved. Extracts published from *Telecomeuropa* with the permission of News Bureau Publications, 3 Princes Buildings, George St., Bath BA1 2ED, UK. Reproduction in any manner, in any language, in whole or in part without the prior written permission of News Bureau Publications is expressly prohibited.

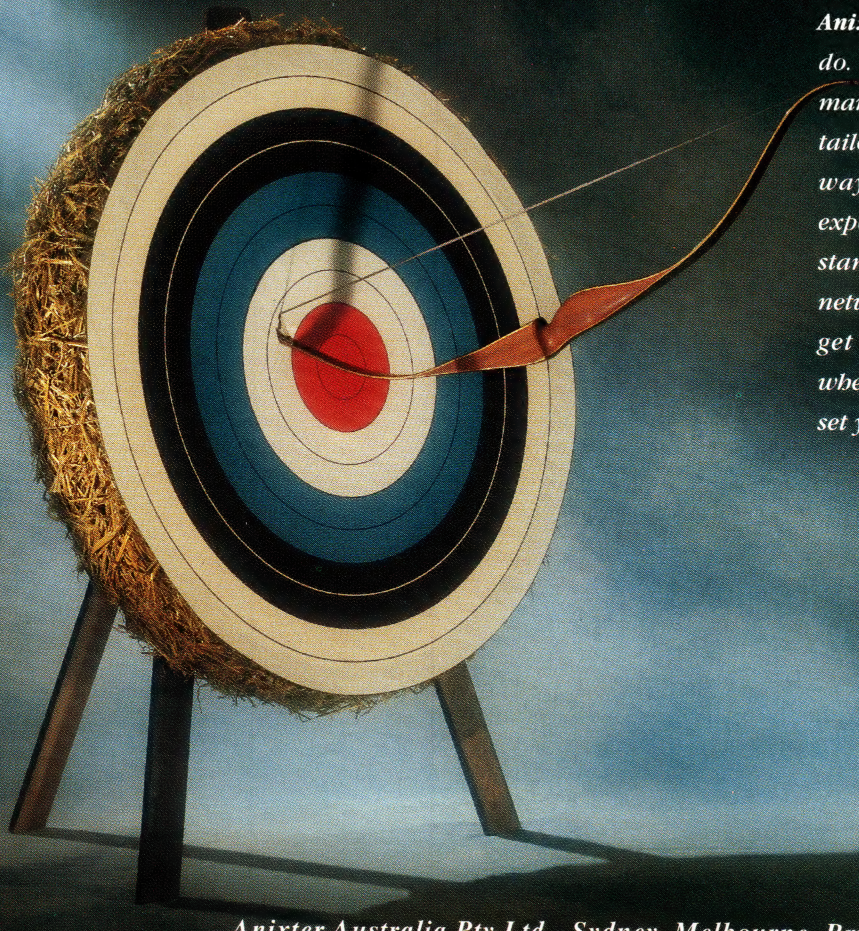

The solution is to work with a company that knows how to take your equipment and make it work for you. A company that knows what's ahead so you won't get left behind. A company that understands how to build a network with the latest routers, hubs and structured cabling systems that will grow with you. Today, the question isn't just about equipment. But who can you trust with the know-how to bring everything in your network together?

HAVING THE EQUIPMENT IS ONE THING.

KNOWING HOW TO USE IT IS ANOTHER.

Anixter. *Nobody knows networking like we do. It starts by working with the leading manufacturers to design a network perfectly tailored for you. Making sure it stays that way. And supporting it with technical expertise from people who really understand your network. For cabling systems, networking products and the know-how to get you exactly what you need, exactly when you need it, there's only one place to set your sights:*

61-2-333-0800

ANIXTER®

Anixter Australia Pty Ltd - Sydney, Melbourne, Brisbane, Canberra

© 1994, Anixter Inc.

Balance Needed in Internet/GOSIP Debate

Alan Lloyd has distinguished himself by his long and concerted involvement in OSI standardisation, both within Australia and internationally. His contributions to the OSI management standards, in particular, are widely recognised. It is therefore no surprise that he staunchly advocates and defends the OSI protocol suite, as his ongoing contributions to *Australian Communications* attest.

Unfortunately, at times it appears that Alan's natural protocol predilection leads to certain exaggerated assertions and improbable comparisons regarding the relative merits of the OSI and Internet protocol suites. I would like to address some of the specific distortions that have appeared in his April and May columns, but first make a few comments about the document that has triggered the present controversy.

The FIRP Report

Alan seems to take umbrage at a recommendation in the Draft Report of the NIST-sponsored Federal Internetworking Requirements Panel (FIRP) to include international 'voluntary' standards such as the Internet Protocol Suite (IPS) within the US GOSIP. Now this particular proposal should be seen in the context of the complete recommendation, which also proposes that federal internetworking be based on a three-level hierarchy, *viz* international voluntary standards (including ISO and IETF standards), then national or consortia standards, and finally proprietary standards with 'multinational commercial prevalence.'

One can debate the wisdom of admitting proprietary standards and the meaning of the term 'international standard,' but it is clear that the Panel sees *internationally accepted* standards as having primary authority in federal acquisitions, and that the Internet suite is not being thrust onto agencies that have made a commitment to OSI or other protocols. Indeed, one of the key concepts recognised in the Draft Report is that of an *affinity group*, a set of agencies or functional interest groups within the government that share information electronically and have common IT requirements. The Draft Report sees interoperability as a staged affair, starting with relative autonomy within affinity groups and proceeding to coordinated approaches across those groups.

More significantly, perhaps, one should note that the recommendation under discussion is the last (and the only technical) recommendation in a set of five. The other recommendations actually propose stronger, more formal, and more pervasive mechanisms for ensuring integration of federal internetworking activities, imposing appropriate responsibilities on the Office of Management and Budget, the Department of Commerce, and the Information Infrastructure Task Force.

I find it difficult to construe the Draft Report as amounting to a mitigation of government efforts to attain interoperability nationally and internationally. It is unfortunate that Alan Lloyd seems to regard even the option of selecting Internet protocols as an attack on OSI, as an attempt by the Internet community to supplant the process and jurisdiction of the international standards bodies.

The fact is that the FIRP proposal denies neither the merits of OSI nor the weaknesses of TCP/IP, but reaffirms the goal of universal internetworking and the need for a single, converged standard as the means to achieve it. The report recognises, quite realistically in my opinion, that neither protocol suite in its current form is likely to provide this global internetworking capability. Of course, the reasons in each case are quite different.

OSI vs TCP/IP: No Clear Case

Most people in the industry will concede several aspects in which the OSI suite is superior to the Internet suite. Such aspects are generally acknowledged as including an effectively infinite network address space, a comprehensive and robust framework for systems management, and a set of functionally superior application-layer standards. These advantages are widely though not universally recognised, and each has a cost: 20-byte rather than 4-byte addresses, no raw connectivity testing capability (à la ping), and application-layer complexity.

The Internet addressing problem needs to be acknowledged but kept in perspective. For example, the current collection of Internet hosts are using fewer than 1% of the theoretically possible four billion addresses. Techniques such as Classless Inter-Domain Routing have been defined in order to prolong the longevity of the existing addressing structure until a long-term solution is forthcoming from the current evaluation process.

Further, it should be noted that the advantages of OSI can potentially be exploited without a total migration to the OSI seven-layer stack. For example, the addressing problem could be overcome by replacing IP with its otherwise similar OSI counterpart, CLNP. And useful OSI applications such as X.400, X.500, and CMIP can be and have been deployed over the Internet and indeed over other kinds of networks, using hybrid stacks such as those defined by RFC-1006 and IBM's Multiprotocol Networking Blueprint. Thus it is not necessary to install an OSI infrastructure to derive the benefits of OSI applications.

OSI has other notable drawbacks — for example, several different NSAP address formats (some are subnetwork-dependent and some are not), and two mutually incompatible network-layer services (connectionless and connection-mode), with no OSI standard for interoperating between the two. And the problem is not fixed by the various GOSIPs of the world, which vary in their support of these options.

Even the oft-cited superior functionality of the OSI application-layer standards has to be weighed against relevant commercial considerations. For example, in the realm of client/server computing, OSI products are thin on the ground — there are currently few products available for transaction processing or remote database access. For simple terminal emulation, the OSI VT solution is often impractical, if available. And in regard to FTAM, perhaps the archetypal OSI application, users struggle with its complexity and bewildering array of options, yet it is difficult to envisage anything much different emerging from the formal standards process as we know it.

I am not trying to denigrate the OSI protocols, which were designed for a valid purpose, but merely to recognise certain inherent problems in them that Alan consistently ignores.

It seems inadvisable to insist that agencies deploy significant OSI infrastructures if there aren't the applications available to run on them, and if they are thereby unable to obtain the compensation of replacing an existing protocol stack.

Can One Protocol Dominate?

There are other reasons why it is unlikely that any single protocol stack can dominate global networking. There is a plethora of networking protocols in operational networks today. The proliferation of Novell NetWare in PC/LAN environments, the vast geographical spread of the Internet and its interconnected networks, and the large installed base of SNA, DECnet, and other proprietary protocols, is likely to restrict OSI to a specifically internetworking role (rather than an 'intranetworking' role) for the foreseeable future. The current widespread use of X.400 to interconnect proprietary mail systems is a good example of the protocol 'equilibrium' that OSI can help to foster.

Also, applications are wedded to protocols, at least to upper-layer protocols. The vast investment in applications implies that their supporting protocols will be around for many years to come and must be accommodated in any internetworking strategy. Note that this situation will change as protocol-independent APIs emerge and gain acceptance.

OSI is Not X.25 or 8802-x

I endorse Simon Hackett's comments [see 'TCP/IP Criticism Unjustified' in the April edition] on confusing OSI with X.25 and 8802-x standards. While X.25 and 8802-x

$4,287,026 saved to-date. $190 saved every minute.

Get your share!

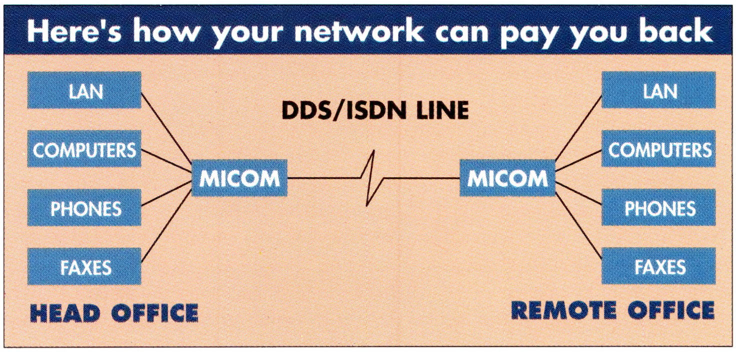

Stop paying out hundreds, even thousands of dollars a month more than you have to on office STD charges. Our MICOM data/voice/LAN network servers send free remote-office voice and fax with computer data over your existing DDS or ISDN line. So why pay for telephone calls and faxes to your interstate offices when you've already paid for a DDS/ISDN line?

With MICOM on your DDS/ISDN line, the 45-minute conference call to your interstate sales office costs the same as the 12-page fax to your local warehouse – NOTHING. And, your computer data and Local Area Network traffic gets there safely, too.

Best of all, a business of any size can afford MICOM. And, it will pay for itself in just a few months.

While you read this, MICOM saved another $190. So get your share of the savings by calling Datacraft, toll-free **1800 803 448** for a free brochure on networking the MICOM way.

MICOM®
Products that pay for themselves.™

DATACRAFT

DATA3472B

Always creating better ways

MELBOURNE	SYDNEY	CANBERRA	BRISBANE	ADELAIDE	PERTH
690 5300	367 4444	239 6811	371 7655	410 0944	311 7111

are OSI subnetwork standards, they don't equate to OSI. The proliferation of subnetworks conforming to these standards predates and is not attributable to OSI. In fact, the dependency is the other way around: OSI incorporates these standards because it was seen necessary to support OSI over these kinds of subnetworks. Note also that the ITU-T X.25 DTE-DCE Recommendation is not an OSI standard — additional ISO standards have been defined to support OSI over X.25 networks. It is misleading to talk about OSI and subnetworking synonymously. When people talk about OSI protocols, they are invariably referring to the top four and a half (or so) layers of the stack, not to the subnetworking layers, which are generic to all stacks.

Also, even within OSI, the connection-mode service offered by X.25 is not the basis on which global connectivity and routing will be implemented. Rather, these capabilities are being delivered using the OSI connectionless protocols, *viz* CLNP, ES-IS, IS-IS, and IDRP.

X.25 and 8802-x support arbitrary subnetwork-independent protocols (e.g. IP and SNA, as well as OSI), and their widespread use does not indicate a correspondingly widespread commitment to OSI, Internet, or indeed any other protocols. In particular, carrier networks do not equate to 'GOSIP' networks!

In regard to the support of the Internet by OSI, it is, of course, true that the carriers support the Internet (by definition, since long-distance connections are provided by carriers), but generally not with X.25. On the other hand, the OSI community actively supports the Internet suite, through its ubiquitous use of the Internet e-mail service. Perhaps more ironically, if one probes inside those 'OSI-compliant' X.25 carrier networks, one is very likely to find TCP/IP running between the X.25 switches!

X.400 is another (useful) component of OSI which of itself does not provide a general-purpose infrastructure for global internetworking. The benefits of X.400 are those of a highly functional set of global messaging capabilities, based on a store-and-forward mechanism (unlike SMTP, which requires a TCP connection between source and target systems). The underlying transport used by X.400 is not important (to X.400), and its addressing mechanism — user-unfriendly though it is — has been superimposed on numerous private non-OSI networks that have X.400 gateways. This is not essentially different from the widespread use of the Domain Name System (DNS) to provide e-mail addresses for many users of private networks that are attached to the Internet.

Telephone System is Not OSI

As an exemplar of OSI, the telephone system is a singularly inapt choice. For a start, the global telephone system was not rolled out according to the kind of grand vision that characterises OSI, but as a piecemeal bottom-up approach to integrating a slew of incompatible and unrelated systems. The global numbering scheme has been retrofitted to the existing network. It was not part of any original blueprint, because there was no blueprint. And the current numbering scheme, with its variable-length country codes, area codes, and local numbers, actually displays one of the acknowledged deficiencies in the IP addressing structure!

Telephone networks, like most carrier networks, have traditionally been 'standard' only at their external interfaces. Internally they use a variety of proprietary or incompatible protocols. For example, analogue signalling protocols are different in nearly every country. PABXs that work in one country will often not work in another. There are no agreed standards for ISDN switching equipment, and implementations of ISDN vary from country to country. Recent standards such as SDH and ATM are virtually the first instances of the adoption of standards for the internals of carrier networks.

I am amused by Alan's use of the term 'OSI-compliant' to describe the global telephone system!

Cooperation Needed

The way ahead for global internetworking seems unclear, but what does seem clear is that neither OSI nor the Internet suite in their current forms will provide the single solution that was originally sought. As Simon Hackett notes in his letter in the April issue, the 'winner' in this struggle will be determined on technical rather than political or other grounds.

It is also clear that the direction charted by the US Government will be relevant to all other countries, and to some extent will limit any discretion on the part of other governments. Certainly, users will increasingly need access to the Internet, not just for global e-mail, but to tap into its vast resources. Any internetworking solution that precluded such access would scarcely be appropriate. On the other hand, OSI will remain part of the US Government strategy — there is no suggestion to the contrary in the Draft FIRP Report. The Internet and OSI communities each have something to contribute in formulating the optimal technical approach to internetworking, and the recent establishment of ISO Category A Liaison status (the highest level possible) for the IETF bodes well for such cooperation.

The key word seems to be *convergence*. There is now good reason to hope and expect that a converged network-layer protocol will emerge from the deliberations between ISO and the IETF, which will hopefully end the politicking on both sides and provide a renewed boost to global internetworking. The Australian Government and other Australian internetworking users should encourage these collaborative efforts between the Internet and OSI communities. It is incongruous to suggest that the Internet can be ignored in the quest for a global solution, or to dismiss it, as Alan Lloyd does, as a useful 'experiment' in internetworking!

In general, the Draft FIRP Report appears to be a basically sound and measured approach to the current problem by a panel comprising long-time advocates of open interoperability within the US Federal Government sphere, and including several OSI protagonists. The following extract from the executive summary of the Draft Report contains a succinct and compelling summary of the critical issues:

'Both the IPS and OSI have strengths and weaknesses, as do proprietary protocols. While a single standard would be preferable, the reality is that there are multiple solutions in networking as in other areas of information technology. Each community within the government should pursue the solution to meeting their mission requirements as a primary goal without a specific technical solution being imposed on them. In addition, they need to consider interoperability with other agencies and the public, existing infrastructure, and cost. These goals are facilitated by standards, but standards alone are not sufficient. The Federal Government needs to be supported by available products and infrastructure.'

Gerard Joseph
Senior Systems Engineer
IBM

Conext Connections

In the May edition you correctly reported Sydney-based Vortex Technologies as the supplier of Scientific Atlanta's Conext videoconferencing system (see 'PCs Bring Low Cost Videoconferencing' starting on page 38). However, further information reported may leave your readers with some misconceptions about our system.

Conext operates on PC ISDN platforms. It is fully standards compliant including H.320 and has been comprehensively tested by Telstra for interoperability with most major vendors. The product operates at speeds from 56Kbps up to 2Mbps (in multiples of 56 or 64Kbps). Though the system has been designed as a desktop unit it is capable of being expanded into a roll-about conference room environment.

A Conext system can be seen at Australia's first retail application for videoconferencing at Citibank's Chatswood Mall (Sydney) branch which is operating at 384-Kbps. This banking application is the most advanced of its kind in the world.

Grahame Bateman
Managing Director
Vortex Technologies

Looking into the Future...

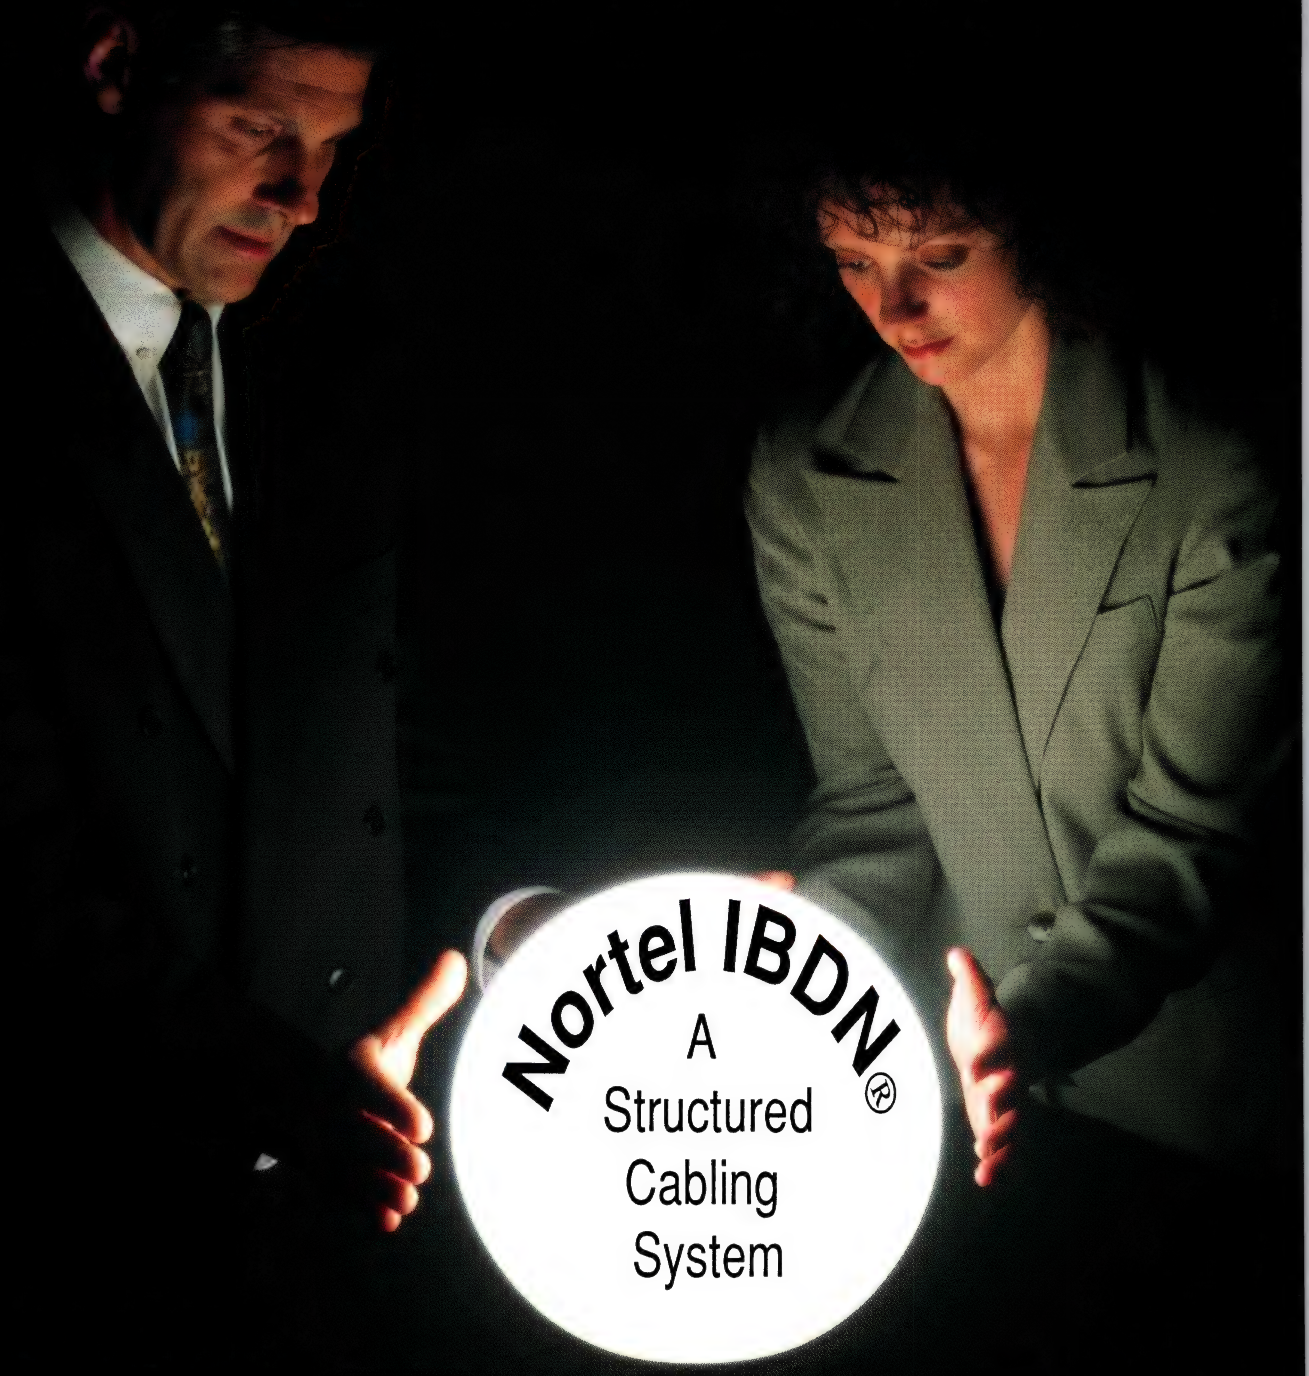

Available Today for the Technology of Tomorrow & Beyond

• Manufactured by Northern Telecom • Supports 100Mbps and beyond
• Integrated Voice, Video & Data • UL Certified Category 5 components • Austel approved components
• Lifetime System Application Guarantee • 15 Year Warranty on Passive Components

For more information on IBDN ®, or a list of Authorised Certified System Vendors and Resellers
Call CableNet, the Authorised National Nortel Distributor

CableNet
6/30 LEIGHTON PLACE, ASQUITH 2077
TEL: (02) 482 7300 • FAX: (02) 482 7311

nt nortel

Telecom Announces Second-Round CATV Winners

Digital Equipment Corporation and Scientific Atlanta are the second-round winners in Telecom's $710 million Cable TV (CATV) project, which will provide Pay TV and other broadband services to an estimated 1.1 million customers in Sydney, Melbourne, Brisbane and the Gold Coast by the end of 1996.

Under the agreements, Digital will provide its Wizard Subscriber Management System, which controls channels and stores customer-related information, including billing information, while Scientific Atlanta will supply the set-top units subscribers will need to decode the cable signal for reception on standard television sets. The two contracts are estimated to be worth around $90 million.

Scientific Atlanta's contract is 'open' and will be reviewed at the end of the year. Gerry Moriarty, Managing Director, Network Products said this is because of the pace of set-top technology change. "Telecom has decided its best strategy is to select an initial supplier so that we can reassess the developments in set-top units early in 1995."

The contracts will be awarded through Philips, which was appointed prime contractor and systems integrator for the project last month in a deal worth $160 million.

The announcement killed industry rumours that Telecom planned to sew up a set-top technology and Pay TV service deal with Rupert Murdoch's News Datacom. The rumours had led to threats of legal action from several firms involved in the original tender process, because News Datacom had not been an original tenderer.

Ericsson Wins $50M Telecom Deal

Ericsson Australia has secured a $50 million contract to provide Telecom Australia with a new network management system. Known as TMOS, the new system will cover the entire Telecom network of around 9 million lines by 1998.

According to Gerry Moriarty, Telecom's Managing Director, Network Products, the new system was selected as part of Telecom's Future Mode of Operation (FMO) plans. Its Australian implementation would require modifications by the Ericsson-Hewlett-Packard joint venture and result in some local software development activity, said Ericsson's Managing Director, Kjell Sorme.

The Ericsson-HP TMOS system is already in use in over 20 countries worldwide, including Sweden, Norway, Holland, Germany and the US. Australian deliveries are scheduled to begin later this year.

Carriers Combine For Asian Link

Optus and Telecom are to undertake a joint project to build an optical fibre submarine cable link between Australia and Indonesia. The 2,00km cable — which is to be known as APCN-A — will be installed in a 'co-operative venture' with Indonesia's PT Indosat and give Australia access to the proposed Asia-Pacific Cable Network (APCN).

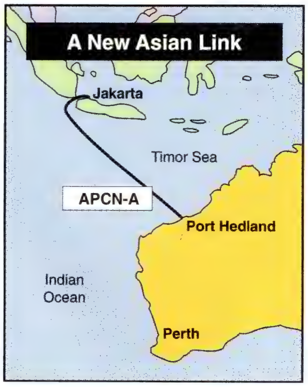

Designed to eventually replace the existing AIS (Australia/Indonesia/Singapore) analogue coax cable, APCN-A will land near Port Hedland in Western Australia. It will use the latest generation optical amplifier and synchronous digital hierarchy technology, and is due to go into service by the end of 1996.

The total cost of the project is estimated at $140 million, including cable supply and laying, and tenders for the contract are expected to be called in July. Telecom will be responsible for 69% of the total investment in the project, with Optus making up the remaining 31%. PT Indosat's contribution will be to take responsibility for the APCN interconnection in Jakarta.

Business Should Tap Into Internet

Access to the Internet should be available to users in government, business and schools, according to a draft report released by the Australian Science and Technology Council (ASTEC). The ASTEC report is part of a study into Australia's future research and information requirements. Internet access in Australia is currently restricted to universities, research organisations or organisations deemed to have 'compatible interests.'

The report finds that wider access to the Internet (which operates in Australia under the name Australian Research and Academic Network [AARNet]) would provide significant benefits to Australia, and that public and private use should be promoted by the Commonwealth Government, which could also use it to disseminate information.

It also recommends that a market-based approach should be adopted for the provision of research data network infrastructure and services.

The cost of operating the Australian end of the network was $4 million in 1993, and is expected to double this year. AARNet is presently considering charging for access on a volume basis, an approach which is endorsed by the report which says the network's charging arrangement should aim at achieving self-funding.

Motorola to Open R&D Centre

Motorola has announced plans to establish a $112 million R&D centre for wireless communications in Sydney, part of a broader strategic investment plan that will see the company spend $400 million in Australia during the next seven years.

The Sydney facility is expected to create 125 new jobs in its first five years of operation, expanding to up to 200 in the long term. It forms part of the Motorola Technology Centre initiative, comprising centres of excellence for wireless communications research, software development and systems integration.

The company plans a Software Development Centre in Adelaide, and has already established a Design and Systems Integration Centre in Melbourne.

The winning of a contract last year to supply a mobile radio network to the NSW Government is said to have been a factor in Motorola's decision to base the new wireless R&D facility in Sydney.

The company has recently signed a Partnership for Development deal with the Commonwealth Government to facilitate expansion of its activities into fast-growing Asian markets.

Matrix Secures Hungary Deal

Matrix Europe Ltd, Matrix Telecommunications' European subsidiary, in partnership with Telecom Finland, has won one of two national paging licences in Hungary.

The $18 million contract delivers the joint venture, which will operate under the name Easy-Call Hungary, a 15-year licence to provide ERMES (the European digital paging standard) public telecommunications paging services throughout the country, with an option to extend the licence for a further seven and a half years. The network will be rolled out in three phases, to cover to 86% of Hungary's population by 1998.

Loading cable onboard the cable ship Cable Venture

New Export Deal For Alcatel TCC

Alcatel TCC continues to win export contracts for its Australian-made submarine fibre optic cables and repeaters. Following its success in March in winning the 3,400km Thailand-Vietnam-Hong Kong (T-V-H) contract, the company has announced a new deal to supply cable and repeaters for the China-Korea Submarine Cable System (known as CKC). The new export success coincided with the commencement last month of cable laying for the 7,080km PacRimWest system.

Valued at $48 million, the CKC contract calls for the supply of 550km of submarine cable and three optical repeaters. The system will link Qingdao, in China's Shandong Province, with Taean, south of South Korea's capital, Seoul and operate at 560Kbps over two fibre pairs. Unlike many other submarine systems, the entire length of CKC will be armoured and buried in the seabed for its entire length to prevent potential damage from the many fishing trawlers operating in the Yellow Sea.

Loaded on board the Cable & Wireless cable ship *Cable Venture* from Alcatel's factory at Port Botany (NSW) last month, PacRimWest will provide a link between Australia and Guam, with onward links to Asia and North America. Its route crosses two of the world's deepest ocean trenches, including the Mariana Trench near Guam, where it will be laid down to a depth of 9,000 metres, making it the deepest submarine cable ever laid. It is expected to be in service by the end of this year.

HP Opens Computing Centre

Lewis E. Platt, Hewlett-Packard Chairman, President and CEO, visited Canberra in late April to open what HP is calling the world's first Enterprise Computing Centre.

A $25 million initiative between HP, the Federal Department of Administrative Services (DAS), and several other software and services companies, the Centre has three major hubs located in Canberra, Melbourne and Sydney. The Canberra hub is located at DAS while the Melbourne and Sydney hubs are located at HP premises. The three sites are linked by ISDN lines.

The Enterprise Computing Centre will be used to demonstrate commercial applications, tools and utilities supplied by its founders on both MVS mainframes and HP-UX systems. Resources available across the three sites include an Amdahl Model 1400A mainframe and several HP minicomputers and workstations. Together, the systems will run 76 commercial software products from 16 software suppliers.

In Brief

ATUG'94 attracted 12,000 visitors according to exhibition organisers, Riddell Exhibition Promotions. The 160 exhibitors covered 13,000sqm according to Riddell's Sales and Marketing Manager, Peter Petherick. ATUG Deputy Executive Director, Richard Allen, said that 1,050 delegates attended this year's conference. ATUG'95 will be held at Sydney's Darling Harbour during the first week of May next year.

Optus says it will invest $683 million in Victoria over the next four years. If achieved, this will bring the new carrier's total investment in the State to $1.2 billion by 1998. The figures were revealed to coincide with the opening of Optus' new International Gateway Switch in the Melbourne suburb of Campbellfield.

Telecom New Zealand is to launch a new subsidiary, Pacific Star Services, to provide a nationwide enhanced facsimile service. The new company will start business within the next few months and create up to 100 new jobs, mainly in Queensland.

Northern Telecom will supply a tactical optical trunking communications system for Australian Army's Parakeet communications project. The system includes ruggedised optical fibre cable, electro-optic converters and multiplexers, can handle transmission of voice, data and image at speeds of up to 4Mbps.

Bull and **TSB International** have announced a strategic alliance designed to provide customers with voice and data network solutions. Bull will now include TSB's telecommunications management offerings in its Network Optimisation Services product.

GEC Marconi Systems has reached agreement with the Defence Science and Technology Organisation (DSTO) to further develop and market high frequency modems. The HF modem technology will provide long distance high speed digital communication over short wave frequencies using error correction techniques developed by DSTO scientists.

MDS (Multipoint Distribution Station) licence auctions will take place in July and August, the Communications Minister, Michael Lee has revealed. The auctions will be managed by the Spectrum Management Agency and should lead to the allocation of 196 licences in 13 areas, including the capital cities.

AT&T GBCS has won a $2 million contract with Westpac Banking Corporation to supply and install a national network of Definity G3i PABXes and Call Centres. The equipment will be used in Westpac's Telephone Banking Project, which aims to centralise and streamline customer telephone enquiry handling.

Telecom Messagetech, a joint venture between Telstra and Matrix Telecommunications, has entered into a $9.875 million research and development syndicate with the Commonwealth Bank to develop technologies for the Asian market. The goal is to develop an enhanced messaging system in several languages, including Indonesian, Malay and Chinese.

AARNet has selected the University of Queensland, Griffith University and Queensland University of Technology to jointly provide a security advisory service for its users. The new service — AUSCERT, the Australian Computer Emergency Response Team — will respond to computer security threats and incidents within the AARNet community, according to Alan Coulter, Director, Prentice Centre at the University of Queensland.

Newbridge Networks, a leading Canadian developer of connectivity products for local and wide area networks, has selected Sydney as the location for its first Australian office. Dale Barker has been appointed the company's Director, South Pacific.

Ericsson Australia has recently received company-wide accreditation to the ISO 9001 quality standard. Ericsson says that only 12 companies worldwide have achieved this status.

Novell has opened its new Asia Pacific Support Centre (APSC), located in the company's North Sydney premises. The APSC will provide second level technical support to Novell distributors, resellers and major accounts throughout the region.

Infotron Systems has landed a contract worth $5 million to supply a packet switching network for the Civil Aviation Authority's new Advanced Air Traffic system. Infotron was awarded the sub-contract by project manager, Thomson Radar Australia.

In Brief

Vodafone has awarded three new three-year R&D contracts to Australian universities. The recipients and contract values are: the University of South Australia ($600,000); the University of Adelaide ($900,000); and the University of Wollongong, in association with the CSIRO Division of Radiophysics ($900,000).

Kazakhstan's Telstra-supplied satellite earth station was officially opened in April in the presence of Australian Governor-General, Bill Hayden. The 16-metre earth station was designed, supplied and installed by Telstra and will provide 1,500 circuits to 12 international destinations by the end of this year.

Acsys a subsidiary of US-based connectivity vendor AMP Inc., has established a sales and support office in Sydney.

QPSX Communications, Telstra's communications technology R&D subsidiary, will refocus its activities following a review of its operations by the carrier. From now on, greater emphasis will be placed on "harnessing the combined research and development capability within QPSX and Telstra," according to Telstra's Group Managing Director, Network and Technology, Doug Campbell.

Privacy Guidelines relating to voice monitoring by the carriers have been issued by the the Telecommunications Industry Ombudsman. They require: prior consumer consent; a continuing overview and regular and independent auditing; and the inclusion of *pip tones* to indicate that monitoring is in progress.

Anixter Australia will distribute 3Com's range of data networking products to network integrators and resellers following a distribution deal sealed between the two companies last month.

Telecom has signed a Master Licence Agreement with Novell. The agreement affords Telecom a special direct relationship with Novell without sacrificing its existing relationships with its chosen dealers. Telecom will soon start migrating its 1,200 servers to NetWare 3.12, but has yet to commit to moving to NetWare 4.0.

Stanilite Electronics has received its third contract from Argentinian cellular operator CTI for the further provision of Cellswitch communications systems. The $1 million contract is for an additional seven Cellswitch devices to join the 85 already delivered.

Stallion Technologies has received funding of $1.1 million from Austrade's International Trade Enhancement Scheme for its Asian expansion plans. A developer of connectivity products for Intel-based Unix platforms, Stallion plans to spend $3.6 million in developing Asian markets over the next few years.

Nira Australia says that Dubbo (NSW) based Western Power has achieved unexpectedly high cost savings through its Freeset business cordless telephone system. As a result, Western Power has halved its expected payback time based on call costs and secretarial disruptions for the system to less than 18 months.

Senator Chris Schacht, Minister for Small Business, Customs and Construction, has officially launched the Queensland-based Cooperative Research Centre for Distributed Systems Technology (DSTC). The DSTC has raised over $50 million to fund a seven year programme researching distributed computing.

Fisher and Paykel has agreed to distribute and market AT&T GBCS products throughout New Zealand. The agreement covers AT&T's Definity PABX and Partner key telephone systems.

Pacific Star will establish an office in Western Australia following its success in gaining a facilities management contract from that State's Department of Land Administration.

Coded Communications, a US-based supplier of mobile data products has signed a Memorandum of Understanding with Exicom to manufacture and licence its products in Australia. Coded has established its Australian headquarters in Adelaide.

The Australian Army has purchased 500 SunSoft Solaris licences together with 100 Sun Sparcstations and seven Sparcservers as part of its Army Command Support System project. Solaris runs across most computer platforms and the Army expects to have it running on over 5,000 computers in the next five years.

EMAA (Electronic Messaging Association of Australia) will hold a public meeting in Melbourne on June 23 to debate the differences between the OSI and Internet communities. Would-be participants should call Michelle Smith at EMAA on (03) 819 6860.

This is the Outback Calling

Telecom claims that the 2.5 million people in remote areas of Australia now have access to to improved communications with the launch of its new Satcom-M and Radphone Direct Dial services. But while the claim is technically true, the costs associated with the new services will dramatically restrict their usage.

Satcom-M uses Inmarsat's global network of geostationary satellites to provide voice and fax communications across the country and internationally. Terminals are available in two configurations: a briefcase sized unit with a built-in or accompanying antenna powered by either mains or battery; and a larger marine version which has a tracking radome antenna and can provide multiple channels. Terminals are reportedly priced from around $20,000 at the low end and usage rates are $1.60 off-peak and $2.30 peak per minute.

Radphone Direct Dial is a public HF radio service which allows users to directly dial most phones in Australia without operator assistance. Monthly charges range from $45 to $60 on top of an annual fee of $99. Telecom also offers a lower-cost operator connected service and says that all its HF radio services offer free two-way airtime for mobile-to-mobile use.

Remote Banking World First

Citibank is claiming a world first with the recent unveiling of its Personal Banking Centre located in Chatswood on Sydney's North Shore.

Together with touch screen, laser printer and scanning technologies, the new facility provides a videoconferencing link to the bank's Sydney CBD headquarters. The Scientific Atlanta videoconferencing system was installed by Sydney-based Vortex Technologies and operates at 384Kbps over a Telecom line.

From the security of a 'private banking booth' customers can converse face-to-face with bank staff while the scanner and other equipment allows for the transmission of documents — providing a unique full branch banking environment, according to Citibank.

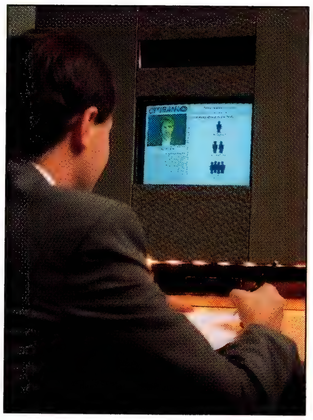

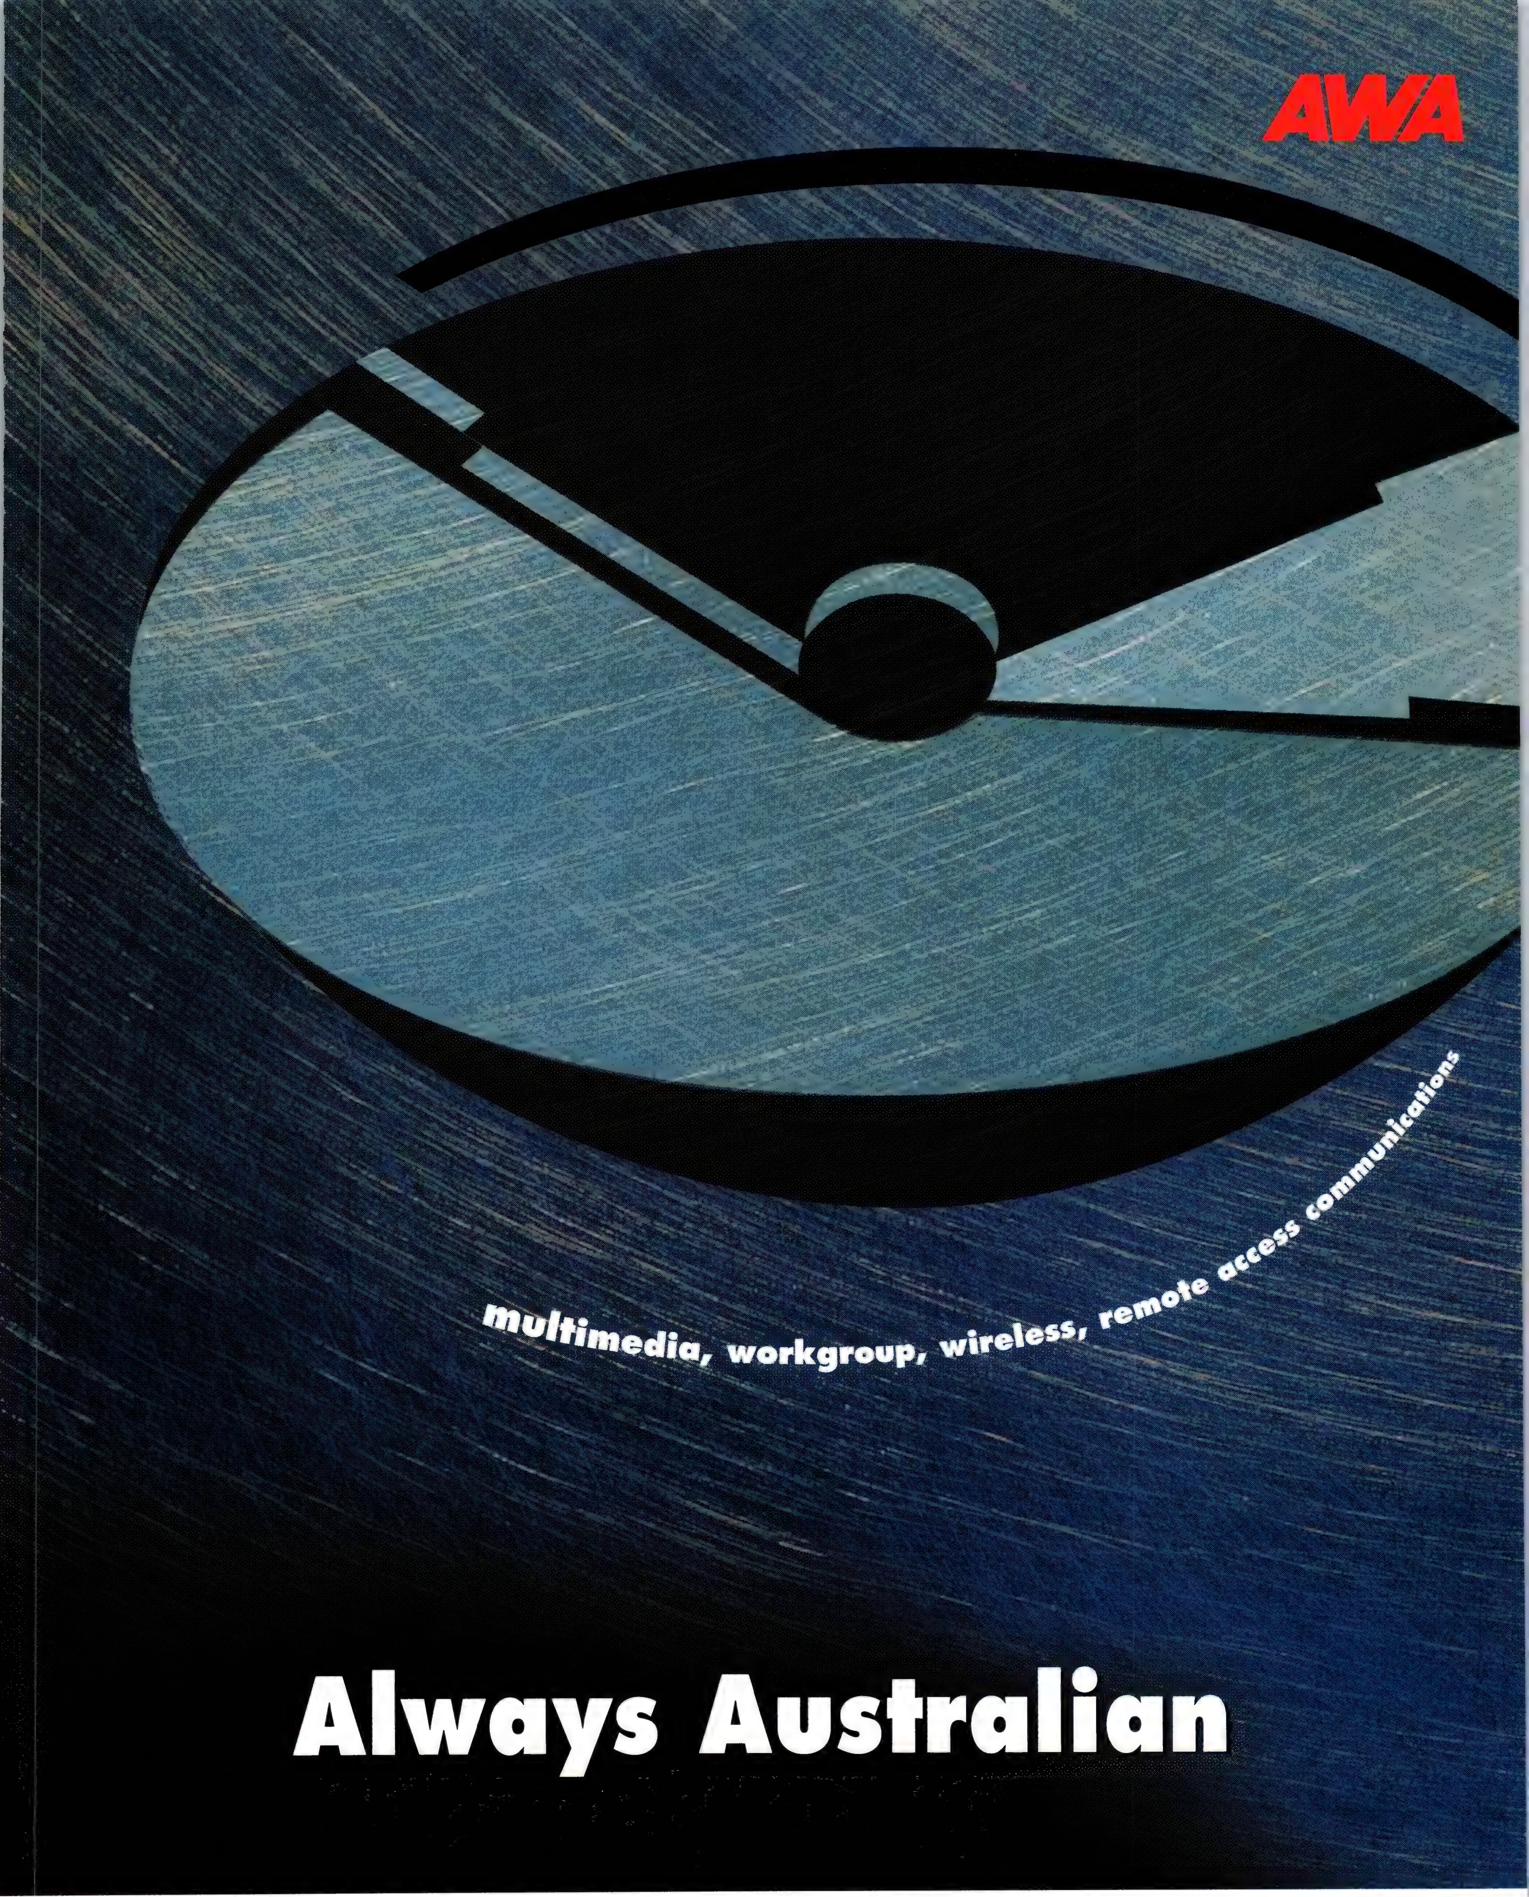
AWA
multimedia, workgroup, wireless, remote access communications
Always Australian

AWA Networks

For more information call
NOW on 1-800-638 9675

NSW
Sydney
Newcastle
Wollongong
Nowra

VIC
Melbourne
Sale

ACT
Canberra

QLD
Brisbane
Townsville
Cairns
Gladstone

SA
Adelaide

WA
Perth
Karratha
Broome

TAS
Hobart

NT
Darwin

In Brief

Scientific-Atlanta has received its first order worth $US15 million for 20 tracking antennae systems from Motorola, as part of the Iridium global communications system.

Novell has appointed a successor to Ray Noorda, the company's long-standing President and CEO. He is Robert J. Frankenberg, former Vice President of Hewlett-Packard and General Manager of that company's Personal Products Information Group. Mr Noorda will stay on as Chairman of Novell's board of directors.

Alcatel Contracting has won a three-year contract worth around $188 million to expand and upgrade Lebanon's telecommunications network. The deal involves the delivery and installation of an optical fibre network based on SDH technology, as well as a microwave transmission network, two satellite earth stations, and a cable network supporting up to 225,000 subscribers.

MCI and Telecom New Zealand have announced the commencement of a virtual private network service. MCI's Vnet international service will be connected to TCNZ's VPN service, providing features such as abbreviated dialling, call re-routing, customised management reports, remote access and call screening.

Microsoft has announced a 30% increase in revenue in its third quarter results, with net earnings for the quarter of $256 million. The company said earnings would have been $336 million if not for a one-time charge of $120 million incurred as part of a lawsuit between itself and STAC Electronics.

Wellfleet has announced its third quarter results for the period ending 31 March, and has reported a 116% increase in revenues to $US103.8 million and net income to $US16 million. The company has also announced a two-for-one stock split.

3Com has reported record sales and orders in its third quarter results, but posted a net loss of $US103.5 million for the period. Officials said this was due to a one-time $US134.5 million pre-tax charge related to the purchase of in-process technologies associated with recent corporate acquisitions.

CompuServe has announced the purchase of more than 100 Windows NT Advanced Server licences for use in the company's New Information Services Architecture (NISA) system, which will be delivered to the company's 1.8 million users worldwide.

PanAmSat has announced Arianespace will launch its PAS-2 satellite in the 191° West longitude this month in French Guiana.

Nokia has signed a contract with Norwegian Telecom for the provision of a Public Access Mobile Radio network. Nokia will supply its Actionet system, which is already in operation in over 20 countries worldwide, including Australia.

Intel has announced it intends to build a $US1.3 billion microchip factory in Chandler, Arizona, which will produce advanced versions of the company's most powerful future microprocessors. The plant is expected to become operational in early 1997.

Telecom NZ has achieved an impressive 390% increase in profit to $NZ528 million on revenues of $NZ2.5 billion for its financial year ended March 31. The figure includes a $NZ350 million after-tax allowance for one-time costs associated with the company's restructuring plan. This restructuring resulted in the company shedding some 3,000 jobs over the last 12 months.

Visa and Metricom have announced a co-operative agreement to establish wireless data communications networks for point-of-sale payment processing. Metricom's MicroCellular Data Network will use radio modems to send and receive transaction messages to Visa's global payment network, VisaNet.

AT&T has signed licensing agreements allowing major Japanese consumer electronics companies Sanyo Electric, Sharp, Hitachi and JVC to introduce and promote new products compatible with AT&T's proprietary VideoPhone technology. The company says it also expects Sony and Canon to become licensees.

Ericsson has won orders amounting to SEK600 million from Malaysia's Cellular Communications Network for the expansion of its ART900 cellular communications network. Ericsson has also recently signed a general purchasing agreement with Telekom Malaysia regarding the delivery of Ericsson Transport Network Architecture products valued at over SEK300 million.

AT&T Clinches $US4B Saudi Deal

AT&T has won the hotly-contested contract to expand and modernise Saudi Arabia's telecommunications infrastructure .

The agreement, known as the Telephone Expansion Project-6 (TEP-6) is worth over $US4 billion, and is said to be the largest-ever telecommunications deal awarded outside the US.

It provides for the provision of a fully digital communications network including 1.5 million lines, GSM wireless services, and associated components for switching, transmission, network management, fibre optics, wireless, civil works and training. It will double the capacity of the current network.

AT&T won the contract over rivals Ericsson, Alcatel, Northern Telecom and Siemens in a highly competitive tender which included all associated activities like cabling and satellite communications systems.

Ericsson, in particular, was unhappy with the outcome, believing it had presented the most competitive offer. Speaking after reports that the AT&T deal had been expanded to cover three times as much as was included in the original tender, Ericsson's Bo Landin, Senior VP, Corporate Markets, said "This information is surprising to us and, if true, is not in accordance with normal tender procedure."

NT, Nokia Get China Foothold

Finland's Nokia and Canada's Northern Telecom (NT) have both announced important deals involving the sale of equipment and services to the potentially lucrative Chinese market.

NT has announced that it will invest $US130 million over the next five years to establish large-scale research and development, manufacturing and sales and service facilities in China for its DMS SuperNode switching family, and for semiconductor and other technologies.

The new NT facilities will be located in Beijing, Shanghai, and Shunde (Guangdong province). To date, the company has sold DMS switches to more than 12 Chinese telecommunications operators for use in public networks, with the number of lines installed or on order now standing at more than four million. The company already has a plant in China manufacturing its Magellan DPN-100 switch.

Meanwhile, Nokia has signed a contract with the Beijing Telecommunications Administration for the supply of a GSM network, which is scheduled to go in to operation later this year.

The agreement covers provision of a DX200 mobile switching centre, register units, GSM base stations and base station controllers, and a DX200 Operation and Maintenance Centre. The system will initially have a capacity of 10,000 subscribers.

IBM Gets Back on Track

IBM executives are finally smiling again following the release of the company's first-quarter 1994 results. The company reported net earnings of $US392 million ($A546 million), compared with a net loss of $US399 million for the same quarter last year. Revenues were also up 2% to $US13.4 billion.

Chairman and CEO, Louis Gerstner, said at the company's annual general meeting that he expected to report a profit for the 1994 financial year, after three years of heavy losses. Revenues had grown in all parts of the world, he said, while expenses had been significantly cut back. Revenues for the Asian region were $US2.3 billion, up 13% over the first quarter, 1993.

While Gerstner warned that the company still had a long way to go, the US Stock Market was buoyed by the result, which was three times higher than Wall Street predictions.

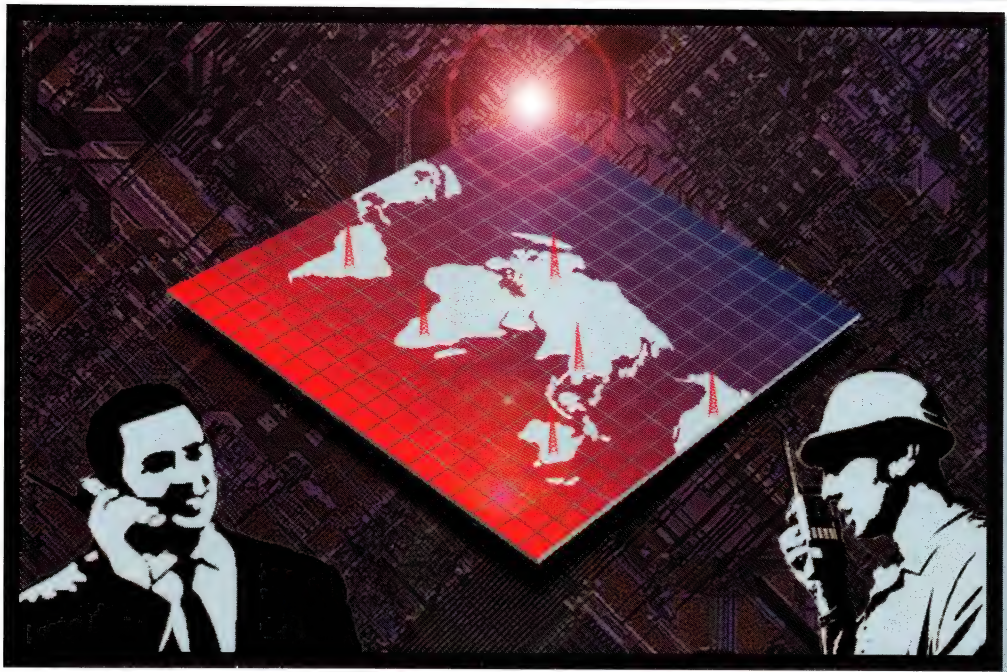

Stanilite
Electronics:
- a major
exporter of
Australian
technology

Stanilite Electronics

- **Cellswitch:**
  Cellular switching systems and networks
- **Trunkswitch:**
  Trunked radio systems and networks
- **Custom PMR systems**
- **Terminal products**

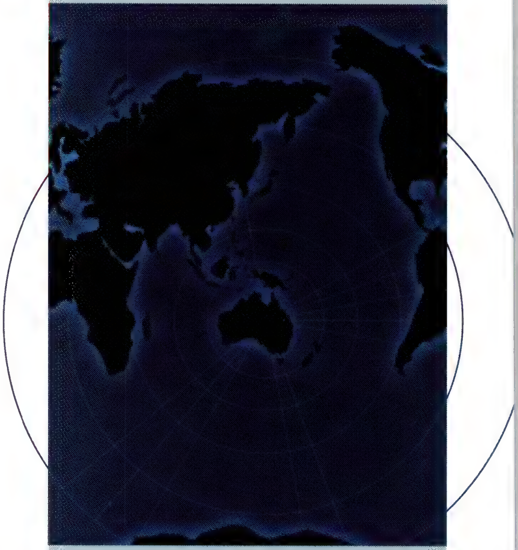

Divisions of Stanilite:

Telecommunications
Products and systems

Projects
Defence & Commercial communications

Technical Services
Customer and project support

Emergency Lighting
Emergency and building services

STANILITE

Stanilite Electronics Pty Ltd ◢ ACN 002 978 403
A subsidiary of Stanilite Pacific Limited
◢ ACN 002 930 289

Head Office
424 Lane Cove Road
North Ryde • NSW 2113
Australia
Phone: (02) 325 1500
Fax: (02) 325 1600

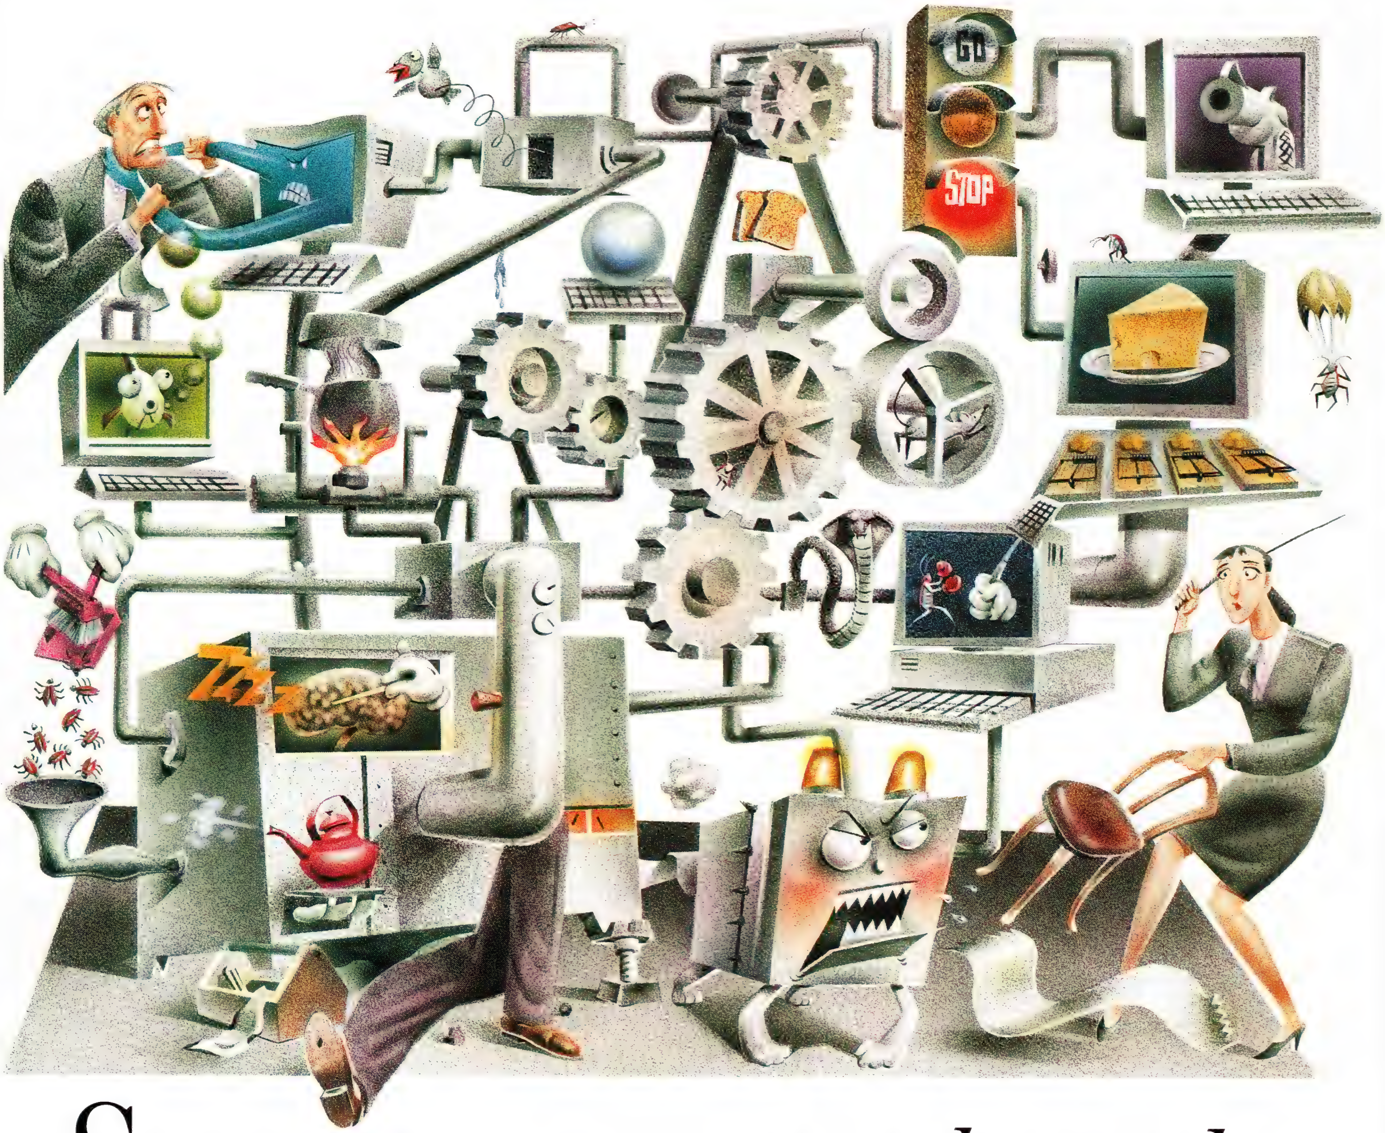

Sure anyone can patch together a computer network... but then...

Complex computer networks needn't be a nightmare - you just need the right partner. Memorex Telex has solved computer networking problems for thousands of companies around the world. We're acknowledged experts in integrating traditional and Open Systems networks within IBM mainframe and AS/400 environments. We have the products, the know-how and a service infrastructure that spans 25 countries.

Let us tame your network!

Call us now on 008 808 680 or Fax us on (02) 805 5758.

Quality System
Quality Endorsed Company

"We connect IBM environments to the rest of the world."

MEMOREX TELEX

3 Thomas Holt Drive, North Ryde NSW 2113, Australia.

IBM and AS/400 are the registered trademarks of International Business Machines Corporation. Memorex Telex and Dynamic M design are registered trademarks of Memorex Telex N.V.

SASS 906ACMEM

Getting a Pay TV Toehold

Philips has won the plum role in Telecom's $710 million cable Pay TV plans, but did the Dutch giant have to underbid to do so?

Just how dramatically Telecom's relationships with its suppliers have changed — and just how committed the corporation is to squeezing the best deal out of them under the policy of 'world best practice at world best price' — has been amply demonstrated in the recently-announced cable and television (CATV) contract with Philips.

Philips Public Telecommunications Systems (PPTS) celebrated a seemingly stunning win as prime contractor and systems integrator — Philips' biggest contract since starting up in Australia in 1926. Newspaper headlines trumpeted the enormous project budget of $710 million over three years; the bubbly flowed at PPTS headquarters at Burwood in Melbourne. But a closer examination reveals just how tightly — financially and otherwise — Telecom plans to hold the reins and control every step in the construction and development of the CATV network; how far ahead Telecom is peering into the future; and how low PPTS was prepared to pitch its bid to get a foot in the Pay TV door. Whatever anyone says, in the end, price was King.

One of the smaller of the 14 Philips divisions in Australia, which collectively generate annual revenues of around $650 million, PPTS has landed an initial contract for $160 million over three years, while the possibility of a contract to supply set-top units cannot be ruled out. With Telecom support, the company is evaluating four different set-top units, two of them its own products. As systems integrator, PPTS will have more than its fair say on which ones are chosen over the next few years.

Despite industry sniping about how much PPTS actually stands to make from its CATV contract(s) (some sources have described its bid as "almost sacrificial") the real victory for the Dutch multinational is that it has secured entry to the budding Pay TV business in Australia. The deal builds on PPTS's existing contracts to supply high-capacity rural radio systems to Telecom, as well as its ISDN terminal adaptor business.

In the US, the company already has a similar CATV deal with Bell Atlantic for a prototype digital television and information network, due to come on-line next year.

Under its CATV contract, PPTS will supply all the telecommunications equipment, including 10,000km of cable from MM Cables, its own cable transmission equipment and a raft of peripherals, such as connectors, from a mixture of US and Australian companies. Ironically, Philips supplied the US-based analogue cable transmission gear used in the Centennial Park (Sydney) CATV trials, run by rival, Pacific Broadband Networks (PBN).

In the broader scheme of things, Philips' slice of the CATV pie pales by comparison with the remaining $450 million of the total budget, the bulk of which will be spent by Telecom's construction wing, Visionstream, whose role will be to install the CATV network's cable backbone, starting this month. Some of Visionstream's expenditure will be recycled back through Telecom as the construction arm leases Telecom-owned sites and office infrastructure.

Formed earlier this year, Visionstream plans to lay cable using existing pits and ducts where possible, to 30,000 allotments or 150,000 homes in its first year, and to 1.1 million homes — a quarter of the national total — in Sydney, Melbourne and the Brisbane-Gold Coast area by the end of 1996. The CATV network will bring optical fibre to within half a kilometre of homes, with coax used for the final link to cut down the expense. Visionstream expects its staff to rise from the present 50-odd to over 1,000 within a year. "By the end of 1996, we'll be connecting homes at an annual rate of 500,000 and we'll continue at that level as we move on to the other capital cities and major regional centres like Newcastle and Geelong," claims Visionstream Managing Director, Bob De Boer.

PPTS General Manager, Brian Adams, told *Australian Communications* that while the short-term application of the CATV network would be a Pay TV service to match Optus's satellite offerings, the difference between the Centennial Park and Wollongong pilots and the final network would be the addition of some two-way, duplex exchange capabilities. "This will position Telecom a bit ahead of the world, particularly the US, in terms of cable TV, and we are offering cable technology that will be critical to the future of interactive services," he said.

Playing down PPTS's consortium links with the cashed-up Lend Lease, computer giant, Digital Equipment Corporation and microwave specialist, Comsys, Adams said one of the major strengths of the PPTS pitch

COMMENT

The Coming of Coax

Telecom says it is installing a 'hybrid optical-fibre/coaxial network' for Pay TV, but despite the pom-poms and twirls of the spin doctors what we are now getting is the same ancient coaxial cable network that has served America so well over the years. And strange as it may seem for me to be praising Telecom, I'd suggest that this is probably the best broadband technology for Australia in the foreseeable future. The digital-satellite-to-the-home licence conditions placed on the Ucom/New World/Lenfest/Australis/Continental Century/Venture Pay TV group are looking increasingly shaky, and while MDS is a good 'last-mile' technology it lacks coaxial's 'interactivity.'

So now we will have Telecom and Optus feeding signals, via the interexchange network or the satellite, to a large number of independent cable operators who will be taking programming from a couple of major wholesalers. Telecom is proposing a coax network which will radiate out from 'head-ends' (main amplifiers) located in your local telephone exchange. It will do this in almost every large town or suburb because the twisted-pair ducting already follows this radiating pattern. Telecom will feed these head-ends via the excess optical fibre capacity that it has around the country for the inter-exchange and interstate phone network. We (the public) have already paid for this over-engineering of the interexchange network through excessive STD call charges over the years, so now Telecom can afford to give capacity away to the cable companies almost for free.

Sixty channels nationwide at 34Mbps each (say, a thousand ADPCM voice channels) should be a cinch with the 36 fibre (2.4Gbps) cables on the main routes, and where they run into capacity problems they can use MPEG-2 compression and knock it back to, say, 8 or 12Mbps a channel. If you want to know what Telecom will be charging its privileged wholesalers for this service, you won't find a mention in the tariff books (and don't try to calculate what this would cost in normal commercial 'bandwidth' terms; you'll have a heart attack — for 70 channels nationwide it's about the GNP of New Zealand!). But you can always estimate one component of a 'packaged' cost by seeing what the (in this case satellite) competition has to offer.

It's far cheaper to use the satellite when you want TV channels distributed to many head-ends, and Optus has transponders sitting idle at a theoretical price of a few million dollars each, which it will obviously use. But if you ask for competitive quotes from PanAmSat, a TV channel (9MHz) is on sale for $US500,000, and Rimsat quotes $US300K (C-band Australia-wide): so add a bit for the receiving dish, and this is what Telecom has as distribution competition.

Of course it will be packaged together with Telecom's street cabling. In Australia, coax will run down current street ducting, and this is actually cheaper than stringing it on poles as US cable companies are forced to do (they are in competition with the telcos). Optus says it has the right to use Telecom's ducting and exchanges also if it wishes. I hear that it does so wish!

Coax networks have a current limit of 750MHz, but apparently Telecom is dropping back on a eight-year old 550MHz technology using Philips/Magnavox amplifiers and passive elements. This cost-effective, well-tested technology is based on 19mm ('750') trunk cable and 12mm side-street splits in a tree-branch configuration with trunks running down about every sixth main street, and cross splits filling in the rest.

Added to all the UHF/VHF/MDS/SHF radio frequencies available, a bandwidth of 550MHz in coaxial probably provides all we could conceivably need in the next decade or two, especially with the 4:1 compression of MPEG-2 video. We are talking thousands of TV channels here.

The coaxial bandwidth division is also internationally standardised these days. From 0 to 5MHz is usually left vacant because of signal ingress from local sources, and from 5 to 35MHz (or more recently to 45MHz) the bandwidth is reserved for the optional reverse channel (signalling and telephony from the homes back to the head-end). The video channels fill the space from 54MHz to 550MHz in 7MHz chunks.

Amplifiers need to be spaced about every half to three-quarter kilometres along each trunk (depending on the quality of the cable and the load) and these amps come with a plug-in slot for back-channel amplification. Compressed digital video can quite happily coexist with the analogue channels later, so they're not tying themselves down to analogue only.

Telecom has licenced the first 10 channels on its cable (about 70 overall) to programmer Cable TV Services (CTS), and the rest are obviously available to other groups — probably the PMT consortium (or at least the P and the M) later when the rollout gets to the economic phase.

Some of the figures Telecom has circulated about the cable network, however, just don't stand up to examination. Telecom claims, for instance, that it expects to connect one in every three homes passed — but if it does this it will feature in the Guiness Book of Records. These high connection rates are only achieved in areas with extremely poor over-the-air TV reception.

Telecom also puts a figure of $700 million on the initial rollout to bypass 1.1 million homes: but $700 a home or $2,100 per connection is ridiculously high. With such a mature technology and ducted cable, the more realistic figure is probably less than $100 per home passed or, say, $500 per home connected (assuming a more realistic 1-in-6 ratio). It will be lower than this in the first year when they concentrate on the high-rise residential blocks.

In most of the world, cable TV has only been marginally profitable to the carriers — mainly because the rights are franchised by local regulators which often screwed down the companies and limited their profits. On the supply side, the Time-Warners and major program distributors screwed the smaller cable operators also. What caused the recent financial interest in US cable companies was the development of interactive technology — in particular, shared-media telephony systems which can use the 0-to-54MHz return channel split for voice. There's plenty of room here for a few hundred telephone channels plus text and control information, and each trunk reuses this bandwidth, just like a cellular base station.

So coaxial 'video' cable has the potential to compete with the telephone twisted pair in providing both low-cost voice telephony and a whole range of interactive multimedia services. Since coaxial is 'shared' media, each home along the line is always connected, so potentially the cable operator will be able to offer many non-circuit-switched (connectionless) services in a way that the conventional phone carrier can't match.

This threat to the local-loop monopoly has been driving up the share value of the cable companies. The broadband superhighway is quite likely to have more to do with coaxial cable than with switched twisted pair — and in most of the world these two loop technologies will be in competition for telephony and data. In Australia we take a more generous approach to our monopolies; when a new technology threatens, we allow our dominant carrier to own both.

Stewart Fist

was that the company already had a national distribution network in place which had "a couple of million contracts per year with customers." The network included subsidiary, United Customer Service, which handles warranties, after-sales service and customer support for Philips products throughout Australia. "Through this facility and our rental company, Rentlo — and possibly Radio Rentals later on — we can deliver a huge consumer market," he said. "At an international level, our company has a history of analysing trends of technology into the home and assessing what will be successful and what will not. Pay TV is a gamble at this stage and Philips can make useful contributions to Telecom's strategic analysis."

Adams said Telecom would also benefit from Philips' controlling share in audio-video giant, Polygram, and television content providers, such as medical and educational programme makers in the US. These would complement any Telecom arrangements with licensed local Pay TV companies such as Cable Television Services, Paynet, Rowcom, Access and Home Show Cable, and specialist programme providers and production houses.

However, Telecom's National Manager, Pay TV Technologies, John McCarter, said Philips' customer and content strengths were secondary issues at this stage and it was Philips' bid price which had clinched the CATV contract. "In industry briefings, we signalled we were going to be aware of cost issues and at the end of the day, Philips had the lowest cost for a quality product," he said.

Win Some, Lose Some

Bitterly disappointed, both AT&T and PBN put on their bravest faces and spoke optimistically about their future prospects in Australia and offshore. According to PBN Executive Director, Trevor Poussard, there are no sour grapes. "Telecom must do what it thinks is best for its future. In the meantime, we will chase niche players, such as the half-dozen local cable TV licensee companies or any others who want to build a cable network, and vigorously pursue our existing business opportunities in China, where cable has been mandated over satellite delivery."

AT&T's Network Systems Managing Director, John Stefanac, admitted to being "very disappointed." But he said there were still plenty of opportunities in Australia for the company's wide range of services and products, from international telephony services to telecommunications hardware, computers, EDI and e-mail. The company had offered Telecom a first-class end-to-end cable and television solution that was winning billions of dollars in contracts with Regional Bell Operating Companies in the US. "Philips may have Polygram, but we have agreements with the likes of Time-Warner and Paramount to develop interactive programming material in the US, as well as Viacom, TCI and huge games companies like Sega and 3DO."

Service Integration

At Telecom headquarters, after the flurry of the FMO, ADSL and CATV contracts, the focus has moved ahead to the new customer access technologies that will be required for the upgraded core networks three to five years from now, and beyond. With the Broadband Services Expert Group concluding its deliberations on the introduction and impact of broadband services in the Australian marketplace, Telecom is conducting a major research project entitled Customer Access Networks-Future Mode of Operation (CAN-FMO). The acronym is designed to make a clear distinction between the public switched backbone FMO, and future modes of delivering multimedia interactive services to the customer.

Manager of the CAN-FMO study, John Nachef, who has recently visited various RBOCs in the US and studied developments in Japan and Europe, says the CAN-FMO represents a move away from overlay networks towards a single access infrastructure. "The key thing we're examining is how best to integrate all Telecom services, whether narrowband or broadband, through a single access infrastructure," he said. "Ultimately, CAN-FMO will allow the advanced capabilities of the core networks to be extended to the customer." This would help stimulate the development and usage of Telecom's CATV and other broadband services, he said. The end result of Nachef's research will be a recommendation, in the form of a business case based on field trials now under consideration, to Group Managing Director, Network and Technology, Doug Campbell, in about 12 months.

Bernard Levy

Stock Watch

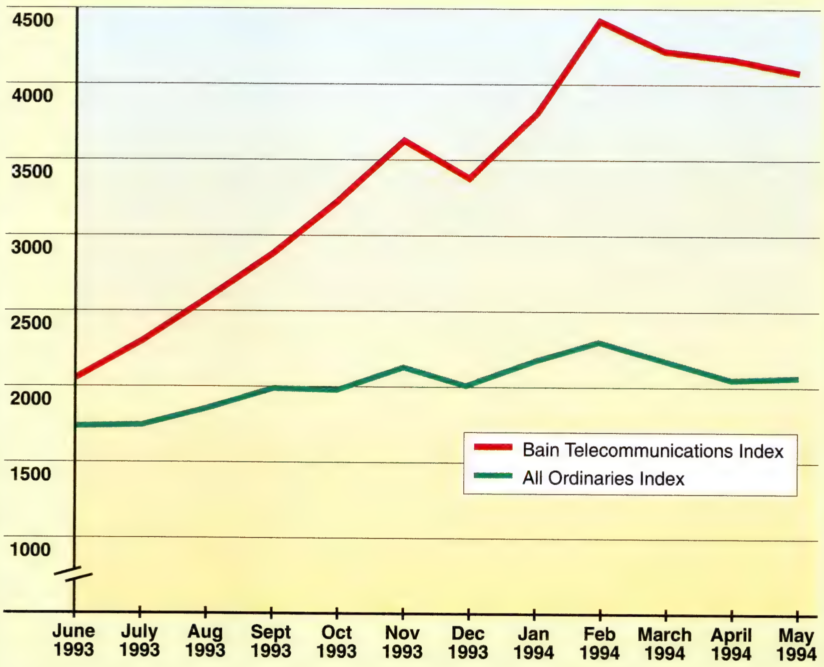

The Bain Telecommunications Index (BTI) includes only those public companies which derive the majority of their earnings from the telecommunications sector. It is calculated in the same fashion as the All Ordinaries Index and is maintained by the Australian Stock Exchange Index Office. The formula is:

$$\begin{array}{c}\text{Today's}\\ \text{Closing BTI}\end{array} = \begin{array}{c}\text{Yesterday's}\\ \text{Closing BTI}\end{array} \times \frac{\text{End of Day AMV}}{\text{Start of Day AMV}}$$

where AMV refers to Aggregate Market Value - i.e. total market capitalisation. The start of day AMV is the previous day's end of day AMV and is adjusted for any changes in the BTI such as additions, capital reconstructions, rights issues etc. This ensures that any movements in the BTI are the result of trading activity in that day only.

Company	Market Capitalisation $M	Performance %	Share Price ($)		
				52 Week	
			1/5/94	High	Low
AWA	189	40.9	0.93	1.02	0.59
Datacraft	90	269.6	0.85	1.13	0.17
ERG Australia	415	157.5	3.20	3.70	1.20
Exicom	96	63.8	0.95	1.40	0.50
JNA Telecommunications	78	39.9	2.63	3.40	1.70
Matrix Telecommunications	131	200	2.70	3.65	1.00
NetComm	8	N/A	0.90	1.40	0.85
Scitec	31	65.9	0.38	0.40	0.07
Stanilite Pacific	176	65.9	5.50	6.70	3.22
Techniche	98	215.6	5.05	6.50	1.49

Market capitalisation is the share price multiplied by total ordinary shares on issue.
Performance is the percentage share price movement over the past 12 months.

Source: Bain & Company

Exports

Telstra Builds on Overseas Success

It may come as some surprise to discover that Telstra now admits to having operations in 25 countries outside Australia. These activities range from nothing more than representative offices to full blown multi-million dollar network investments and are part of Telstra's strategic plan to become a regional, if not a global, carrier with 20% of its revenues generated offshore.

However Telstra's overseas strategy has not been without its critics. A year ago, Terry Cutler, Managing Director of Cutler Consulting Pty Ltd and former Director of Corporate Strategy with Telecom Australia, was quoted in a daily publication associated with the Asia Telecom trade show questioning whether Telstra had a true globalisation strategy. He believed it needed access to far more capital than it had at the time and needed closer ties with key suppliers; long term partnerships which include a fair degree of shared risk. "It will need to form long term partnerships with a significant level of financial commitment and shared risk. So far Telstra has not got beyond expressions of goodwill," he said.

He said many of the multinationals saw these partnerships as important because they were the main means of differentiating themselves from their competitors, and many have given their Australian subsidiaries mandates to develop export business in the region. He also cited Telstra's lack of success in domestic joint ventures and uncertainties about future privatisation as possible handicaps overseas.

Philippines Now, India Soon?

A year later, that uncertainty about privatisation remains but Telstra has massively increased its commitment to overseas investment in its bid to be a regional power, and at least one of its major ventures has gone beyond 'expressions of goodwill.'

In April Telstra announced its biggest ever offshore commitment: $267 million over three years for development of networks in the Philippines through a joint venture company, International Communications Corporation (ICC).

At the time of going to press, Telstra was also believed to be very close to finalising arrangements to make an even bigger investment in India: $280 million over five years. However, there are two major differences between the two projects: in the Philippines the Government's commitment to liberalisation and private sector participation in its telecommunications infrastructure is now clear cut and unambiguous. In India, this is far from the case, and numerous network expansion projects planned with private sector and foreign carrier participation are on hold while the country sorts out its communications policy.

In its Philippines venture Telstra is in a partnership with Benpres Holdings Corporation, owned by media magnate Eugenio Lopez. They jointly own ICC, with Telstra holding 40%. ICC was last year granted a licence to operate an international gateway facility in the Philippines and, as part of its licence to operate this profitable facility, it is required to install a total of 300,000 lines of basic telephone services to help expand the Philippines' under-developed network.

Similar conditions apply to all licenced international and cellular operators, who must install 300,000 and 400,000 lines respectively in a mixture of areas determined by the government such that each operator makes an equivalent contribution to the network expansion.

At present the former government-owned Philippines Long Distance Telephone Company (PLDT) is the dominant long distance carrier with about 90% of the market, but this position is now being challenged. Separate from their local network commitments, ICC and Global Telecom, an international and cellular operator 38% owned by Singapore Telecom, are jointly planning to invest $140 million in a rival long distance trunk network which will link the 11 areas designated by the National Telecommunications Commission where local networks are being installed.

Telstra has been awarded the project management contract for the network, and according to Telstra's Managing Director International, Warren Grace, it will "provide potential export opportunities for the Australian telecommunications industry."

In India things are not moving so fast, however the pro-liberalisation chairman of India's Telecommunications Commission was due to visit Australia in late May for talks with Telstra and the Department of Communications, according to a report in *Telenews Asia*. India has already expressed interest in Australia's approach to telecommunications. Last year also, Indian telecommunications officials met with representatives of the Department of Communications and Austel to discuss liberalisation issues. Telstra has reportedly offered network management services to India's Department of Communications, and has also offered to train 100 Indian engineers as well as seeking partnership opportunities.

Telstra is a member of a consortium granted a licence to operate cellular services in India in 1992. It is still waiting for this opportunity to be fulfilled after the original contests were challenged by one of the unsucc- essful bidders. It has however secured a contact in India, backed by AIDAB funding, to set up a network management centre in Madras for the national public network.

Earlier this year, Telstra's Indian subsidiary, Australian Telstra Holdings, entered a joint venture with Indian Amadeus Finance and Investment Ltd to operate data communications facilities for the new Indian National Stock Exchange, based in Bombay. The new joint venture company also intends to offer value-added services to other Indian companies.

A Chinese Puzzle

In other regional forays into network operations Telstra last year acquired a cellular spectrum licence in New Zealand for $10 million and was part of consortium to win a domestic carrier licence in Hong Kong.

So far there has been little indication of Telstra's plans for its New Zealand spectrum licence. $10 million seems like a lot of money to be lying idle while two competitors are doing their best to sew up a market of only four million people.

In Hong Kong a consortium between Telstra (20%) and Hutchison Telecommunications (80%) was one of three consortia granted licences to compete in fixed domestic networks against Hong Kong Telecom. Observers saw the licences as having little immediate value, there are no local call charges in Hong Kong, but as being sought after as good means of getting into the mainland market after Hong Kong becomes part of China in 1997.

There has been considerable speculation in Hong Kong that Telstra's involvement is far from certain. The *Asian Wall Street Journal* reported that Telstra might opt to be a strategic partner rather than sharing in the estimated $US450 million planned investment. More recent reports suggest that in fact it is Hutchison that wants Telstra out. Telstra's Managing Director Corporate Government and International, Steve Burdon said in May that, in fact, no agreement with Hutchison had yet been signed, but said he expected this to happen "sooner rather than later."

In Saudi Arabia, Telstra has had a multimillion dollar network management contract since 1988. But the Kingdom's recent decision to grant a $US4 billion contract to AT&T to expand and modernise its telecoms infrastructure may signal the early end of Telstra's involvement.

Whatever its real successes overseas, Telstra has taken a noticeably higher profile in recent months. Its sponsorship of The World Infrastructure Forum to be held in Jakarta in October, of an Australia-Indonesia business forum being held in Jakarta in June, and of *The Economist's* Roundtable with the Government of Pakistan in Islamabad which took place early in May, are all helping to get the Orange Blob well known in the region.

Stuart Corner is the Managing Editor of Exchange and Editor of Telenews Asia.

Communicate At Breakthrough Speeds.

Get AT&T's SYSTIMAX® Structured Cabling Systems, And Get Speeds Beyond 100 Mbps.

Over copper. Or fibre. Your network can now have unprecedented speed and interoperability. All by installing AT&T's SYSTIMAX Structured Cabling Systems, with our HIGH-5™ product family of Category 5 components. They support today's network applications and emerging technologies like ATM, up to speeds of 155 Mbps. And this total solution offers you end-to-end connectivity, maximum network compatibility and the ability to evolve your network as your future needs change. All from a single, reliable source, AT&T's HIGH-5 products exceed EIA-568 Category 5 specifications, and every certified installation includes AT&T's 15-year extended product warranty and system assurance.

The copper cabling was created by AT&T Bell Laboratories and has become the benchmark for Category 5 standardisation testing. So find out how AT&T's SYSTIMAX Structured Cabling Systems can help you break through communications barriers. Call us on **008 804 738.**

AT&T SYSTIMAX® Structured Cabling Systems. The last cable system you'll ever need.

AT&T
Network Systems

For more information fax us on (02) 352 9279 or send this coupon to AT&T Australia Limited, Locked Bag No. 19, North Ryde, NSW 2113.

Name ....

Address ....

.................................Postcode ....

AC

© 1993 AT&T ATT4000

Tap into Asia

"Most Australian telecommunications companies believe exports will fuel their revenue growth in 1993-4 year ... particularly in South East Asia and China."

EXCHANGE NEWSLETTER 2/7/93

Don't miss out on these opportunities. Get the latest news on these crucial markets from Telenews Asia, the region's only fortnightly dedicated to telecommunications.

There's no advertising. Just concise, objective and independent information, supplied by our Asia-wide network of correspondents and edited for specialists by specialists.

GET THE NEXT 3 ISSUES FREE

Judge Telenews Asia for yourself. Just mail or fax this form for your 3 free issues.

Please send me the next 3 Telenews issues at no charge and without any obligation.

Name ____________________

Company ____________________

Address ____________________

MAIL TO: STUART CORNER INFORMATION SERVICES
PO BOX 13 ROZELLE, NSW 2039 AUSTRALIA
FAX: (02) 818 2294 INTL: + 61 2 818 2294
PHONE: (02) 555 7377 INTL: + 61 2 555 7377

TELENEWS Asia

Optus

Optus Opens Up

In April, Optus held an 'exclusive' press briefing with a small group of technical journalists for no other reason than to fill in the annual record and make company executives available for questioning. CEO Bob Mansfield and a half-dozen of his top staff fielded question for three hours over coffee and sandwiches.

The journos came away from the meeting all buzzing with excitement — not so much from the information provided, but from the mere fact of the event. After all, there must be something in the carriers' ethics manual that prohibits this — it's certainly never happened in the past!

The first surprise was Bob Mansfield's announcment of a very early date for Optus to break even. 'Early in 1995' is apparently the official word, with the financial year 1995/96 expected to be the first to show the company as profitable. Revenues for the year 1992/3 were about $300 million, and the 1993/4 figures would come in at around $800 million, Mansfield added.

Optus's bottom line is helped to a degree by the fact that the company still owes the government $300 million for its licence, and doesn't begin to pay this money until January 1995 — and then only in $75 million blocks every six months.

Optus also apparently gained favourable finance for its capital investment through the Commonwealth Bank. It sold the bank the Hughes B1 satellite for $300 million, and then leased it back. Some $800 million was also raised through a Westpac syndicate, but recently the coffers were running a bit low, which is why Packer was bought in, according to Mansfield (though some reports suggest a different motive). But Packer's entry into Optus wasn't discussed in detail because the deal had not been finalised. We were assured, however, that both Optus Pty Ltd and Optus Communications will be listed and floated together in 1995.

About half the company's total revenues come from long distance telephony services, and will continue to do so in the immediate future. The remaining 50% is divided between the Business Networking Services (BNS — mainly satellite services) and the resale of Telecom's AMPS mobile cellular. At present, the mobiles division earns four times the income of BNS, but the ratio is expected to shift to about 3:2 as the BNS revenues include more of the Central Business District fibre component.

On the BNS front, Optus claims it does not see PanAmSat (which is due to fly in June) as threatening its satellite business because the Americans are only offering four transponders. All up, the three Optus satellites have 45 continental transponders (although A2 is used entirely for telephone backup).

But this is wishful thinking on the part of Optus: PanAmSat can offer a total of 4 to 7 Ku-band transponders (three are switchable with China) and 16 C-band transponders with a footprint over Australia. What's more, users will be able to up-link to this satellite directly out of the US. Rimsat has also appeared over the Pacific to complicate matters further: in one to two years the Pacific is going to be awash with satellites (see 'The Great Asia Pacific Space Crush,' in the April 1994 issue).

Mobilesat, the L-band all-Australian mobile voice/data system on the new B1 satellite, will be up and running in September 1994, Optus says. The company still can't quote a handset price, but the word is below $7,000.

This year has seen local switches and an early business customer access network in place in Adelaide and Perth, and the company is expanding its fibre-laying in Sydney, Canberra, Melbourne, and Brisbane. Optus's big-city network topology will consist of two rings of fibre, one around the CBD, and a wider ring in the outer suburbs. The outer Sydney ring, for example, will pass through Belrose, Hornsby, Parramatta and Rosebery. Optus is also using microwave point-to-point links in some cities to bring more distant corporate facilities within reach of the CBD or outer fibre-ring.

The company's expansion of its intercapital network is obviously progressing well, with some 5,000 kilometres of cable now in the ground and new trunk switches in Adelaide and Perth; channels are rented from Telecom to fill in the gaps in the map, including the Nullarbor crossing. On the east and south coast, Optus has now reached Brisbane and Adelaide with a full 36-pair fibre cable, and has doubled-up on the Sydney-Melbourne 36-fibre connection with a second cable via Wollongong. This network is totally SDH running at 2.4Gbps, and currently only one pair in each cable has been fired up. For backup and overload, Optus uses the A1 satellite.

GSM is still central to the company's plans for a mobile future. The company believes that all three networks have about 10,000 GSM digital customers, claiming about 3,500 customers for itself; however, GSM still only represents one in every hundred mobile handsets out there.

Optus is differentiating its GSM service with its 'Surepage' Short Messaging Service which, (for an additional $35 a month) provides up to 100 messages, each of up to 160 alphanumeric characters, on the new Nokia handsets. There's also Voicemail with a free message-waiting indicator.

'Metroplan' is the company's new discounted billing system for those GSM users who generally stay within a local area — it's cheaper, but you pay double if you use your handset interstate. Optus also has a GSM rental-phone system in place, and earlybirds can take advantage of a special 'try-free-for-two-days' promotion.

There is no doubt in Optus's collective mind that CDMA and American Digital AMPS (TDMA) are both banned from use in cellular mobile systems in Australia. According to both Ian Boatman, Optus's Chief Operating Officer, and Bob Mansfield, the company's agreement with the government is that only GSM is to be permitted.

However, Ian Slattery of Austel denies this; he claims that GSM is only the 'first' digital system to be licensed, and that there is nothing hindering the introduction of CDMA. Expect a considerable bun-fight if Telecom decides to bring in CDMA technology (which uses the AMPS band).

Dr Wayne Nowland, Optus's Chief of Business Operations, has recently metamorphosed into the guru of 'data superhighways.' He spoke at length on the company's direct-digital satellite Pay TV plans which have been 'temporarily' discarded in favour of MDS delivery over the last-mile. MDS terrestrial transmissions will be supported by satellite distribution to the masthead (digital DBS will come only in mid-95, he says). Both Nowland and Mansfield hinted very strongly that Optus is about to challenge Telecom in the laying of street coaxial, and it certainly makes sense to use the satellite for wholesale distribution only.

Optus executives were also scornful of ADSL (Asymmetrical Digital Subscriber Line, see 'ADSL — Bridging the Superhighway Gap' May 1994 and this issue) as a useful Pay TV technology; they expect coaxial (with some limited MDS) to be the video-last-mile mechanism for the next decade.

The company is holding discussions with a large number of small cable licensees in larger country towns and regions, and is also investigating opportunities in high-density residential areas in the major cities.

Andrew Bailey, who looks after Optus's regulatory affairs and legal issues discussed the question of when Telecom may cease to be the dominant carrier. This question occupies most Optus minds, because when Austel ceases to treat Telecom as 'dominant,' Optus will be forced to renegotiate its interconnect fees — but this time without Austel's intervention.

On the vexing question of Special Pricing Arrangements (SPAs) and Flexi-Plans, Bailey and Manfield are of the opinion that Austel bends too easily under pressure from the dominant carrier.

Bob Mansfield pulled no punches regarding Austel and its chief, Robin Davey. Austel did a good job for the first three years, he said, but in the last 12 to 18 months he claims there has been a very significant deterioration. He believes the Austel Board tends to use black-letter interpretion, rather than seeing it role as promoting competition.

Stewart Fist

Optus

A Warm Welcome For Packer?

What may have looked at first like another smart initiative on the part of the Nine Network to acquire a stake in Optus Communications now appears to have been the result of pressure on the Optus Board from the carrier's controlling overseas shareholders, BellSouth and Cable & Wireless.

It has also now become clear that not all the company's Australian shareholders liked the idea of offering Kerry Packer a stake, with some fighting hard before caving in. A finance industry source close to the negotiations suggested the real impetus for a Packer presence derived from the foreign shareholders' dissatisfaction with the business and political strengths of Optus's existing Australian shareholders, Mayne Nickless, AMP, National Mutual and the Australian Industry Development Corporation (AIDC).

Some Board members are reported to have taken personal affront to this. "With Telstra still wielding so much power and influence, what C&W and BellSouth wanted was more clout in Canberra and greater input into the decision-making processes affecting the future development of telecommunications in Australia up to and beyond 1997," the source said. "Kerry Packer has influence within all political parties, and his international contacts and television programming skills are invaluable."

Significantly, aside from its stake in Optus, C&W is understood to be in the final stages of winding down its direct presence in Australia, having sold off the various divisions of TR Services late last year, almost two years after Optus's bid for the second carrier licence. C&W's nearest regional headquarters are in Hong Kong. BellSouth, meanwhile, continues to maintain a direct local presence via controlling interests in Link Telecommunications and Mobile Data.

In what appears to be a win-win situation for Optus and the Nine Network, the carrier received a welcome cash injection of $318 million to be spent on network roll-out and development and broadband technology, while Packer gained a secure and profitable 15% long-term holding which gives him an alternative avenue into Pay TV, besides his involvement in the Packer-Murdoch-Telecom (PMT) syndicate.

The Nine Network is committed to retaining its stake in Optus for at least five years, and Managing Director of Consolidated Press Holdings, Brian Power, will sit on the Optus Board. According to a senior executive with one of the Optus shareholders, who asked not to be named: "There had been an awareness for some time of the need for some strategic moves on broadcasting and telecommunications convergence. Boston Consulting was brought in and came up with a whole list of names, from Australian to international figures, as potential links in the convergence chain.

"The report was useful but not pivotal. Bob Mansfield spoke to just about everyone in the business around that time, including Packer interests, who had been approached to invest in Optus two and a half years ago, but rejected the idea. But it would be fair to say that in deciding if the Nine Network was the best choice, the issue arose whether the shareholders wanted to experience what happened at Westpac. There was some disagreement, but it would be wrong to imagine that anyone was muscled into accepting the decision.

"It has to be remembered that during the Westpac saga, Al Dunlap headed Australian Consolidated Press and it was a cut-and-burn policy in terms of getting into companies and making them leaner and meaner. But the Packer organisation today is quite different under Brian Powers, and is diversifying into petro-chemicals, casinos, racing, Hollywood movie houses an so on.

"The Packer presence in Optus will probably mean less kicking heads within Optus and more kicking heads outside Optus." The source suggested that while AMP and National Mutual "paid considerable regard to the experience of C&W and BellSouth and believed the overseas groups had the right idea," Mayne Nickless and the AIDC at first resisted the idea of a Packer presence and needed a lot of persuading."

The Packer move into Optus further demonstrates the relentless convergence of ownership of Australia's media and telecommunications operations. Both Rupert Murdoch and Telstra have bought stakes in the Seven Network while still involved with the Nine Network in the PMT syndicate. Telstra's Director of Corporate Affairs, Chris von Willer, while stressing that no single party was authorised to speak for PMT as a whole, said Telstra management was not overly concerned about the Nine stake in Optus, saying it was "business as usual."

"From Telstra's point of view, the PMT Pay TV group continues to be solidly based and there is a good deal of momentum within the consortium towards implementing a viable business plan initially based on the forthcoming MDS licence auction."

According to Hank Spier, General Manager of the Trade Practices Commission, which works closely with the ABA and has special powers under the Broadcasting Services Act, "The whole area of broadcasting and telecommunications is so fluid at the moment, because of convergence, that we are keeping a close eye on all movements. The TPC and/or the ABA might step in if, for example, the Nine Network did a deal whereby it gained some kind of exclusive access to Optus satellite services for its programmes and locked other companies out, which is highly unlikely."

According to Hambros Telecommunications analyst, Dr Chris Beare, the Nine Network has made "a good opportunistic investment. Nine has an agreement that Optus will float on the stockmarket within a year or so, and knowing the interest worldwide in telco stocks, Optus shares would be far higher than the $1.20 per share paid by Nine, even though it is locked in for five years."

Shadow Communications spokesman, Senator Richard Alston, described the Nine Network deal as "healthy" in that it gave the broadcaster a stake in the telecommunications industry and an opportunity to establish a broadly-based Australian multimedia communications company.

"On the negative side, there are cross-media rules in place designed to promote diversity of ownership and maximise competition, but they have been singularly ineffective to date. What's needed is a review of their effectiveness, rather than ripping up the book and starting from scratch."

Bernard Levy

COT Cases

Telecom the Biggest Casualty

It was not surprising to discover some Telecom executives beaming with 'delight' at the Minister's announcement that Robin Davey will not re-apply for the Austel job.

Austel's COT Cases report, which is very much Davey's handiwork — his 'swansong' as one observer put it — is a devastating indictment of the sluggish service culture that has flourished for years inside our national telephone company. Davey says the report was a lesson in 'objectivity' and admits to combing drafts to remove any emotional language.

The overall tenor can be gleaned from the way Austel underscores Telecom's failure to concede that there may be faults in the network, but they can't be found. This attitude, Davey observes, is "perhaps attributable to a hangover of a culture from days gone by of 'never admit to a fault' based on the assumption that if it is denied, it will go away and if it does not, how can the customer prove otherwise?"

The combination of Telecom resistance and COT persistence proved explosive. The regulator initially offered to act as 'an honest broker,' but came up against what Davey refers to as the internal Telecom 'loop': Austel would escalate issues to the top and then find these referred down the line to lower level managers in the very areas that had failed to handle the original complaints.

With media coverage turning up new COTs and the prospect of this 'loop' continuing indefinitely, Austel decided to direct Telecom last August to undertake service monitoring and testing and to supply documentation on eight customers.

The national phone company finally bowed to the inevitable and moved into damage control, commissioning an audit of internal procedures from Coopers and Lybrand and a network review from Bell Canada International. According to Davey, the turning point came when Ian Campbell, Managing Director, Commercial Business, was given control. Campbell conceded that the network was subject to faults; that recurring problems needed to be fixed; that Telecom's handling of customers was not perfect; and that staff may have made inaccurate or rude statements such as: 'You are the only one in the area with the problem' or 'Telecom has no liability.'

The company has so far paid out around $900,000 in compensation.

No Ring Received

Austel's investigation focused on eight business customers who had experienced similar, intermittent and long-standing faults such as: no ring received, busy when not, call drop out and recorded voice announcements when not disconnected.

Telecom refers to these as Difficult Network Faults (DNFs) because they can't be isolated and fixed to the satisfaction of customers using routine monitoring/testing procedures. The use of DNF unfortunately makes it easy to speak about 'difficult' customers, a marginalising label no-one wants.

Each of Austel's eight COTs said they had 'frequently' experienced (often with other faults) the incoming call fault known as 'no ring received' i.e. a caller dials their number, hears the ringing tone, but no ring is received at the customer's premises. Neither the customer nor the caller identifies this unless they compare experiences.

As Austel observes, Telecom's stance is that fault reports are customer perceptions and not necessarily 'genuine faults.' According to the company's Director, Regulatory: 'It is only when our technicians can trace and identify the cause of the perceived fault that it is possible to confirm the occurrence of the reported fault.'

The dilemma for customers seeking compensation for loss of business because of 'perceived' intermittent NRR faults has been finding the 'hard evidence' required by Telecom to prove these occur with a certain frequency. But as both Austel and Coopers and Lybrand discovered, even Telecom itself could not provide this 'hard evidence.' After detailed investigation of Telecom's systems and procedures, both studies found that fault reporting/recording was inadequate and incomplete and that the monitoring/testing equipment was unsuitable and unreliable. Yet customers had been told that 'no fault could be found' rather than 'there may be faults, but we can't find them.'

Telecom does not therefore know the extent of this specific DNF on a national basis. The Telecom/Austel survey of 2,644 small business customers in nine exchange areas by undertaken by Reark Research last November found 105 customers claiming one or more of these problems had affected their business 'significantly' or 'very significantly' during the past month. Just over 50% of this group claimed a 'no ring received' fault (2% overall). The survey did not identify however whether this was related to CPE, the network or both.

Customer Charter

At the press conference held to respond to Austel's report, Telecom CEO, Frank Blount, said the company 'substantially' agreed with most of Austel's recommended changes and had moved already to implement Coopers and Lybrand's recommendations.

Steve Black, the Group General Manager, Customer Affairs appointed in December to oversee the changes, is moving on all fronts. An internal action list reveals new fault recording and management procedures; revised voice monitoring guidelines; mediation/conciliation and arbitration processes; a national customer response unit and a new complaints handling pilot in June, with progressive national rollout from July. Meanwhile, the decision last year to accelerate the replacement of ageing analogue exchange equipment is expected to address many recurring faults (see *Interview* on page 51.)

The change which should benefit all customers is the promise of a Customer Charter which, in simple language, will spell out customer rights and Telecom obligations in relation to the standard telephone service. The pamphlet will cover areas such as as making/receiving calls, entry in the White Pages, copies of White and Yellow Pages and use of fault reporting and repair services, as well as a list of useful contact numbers for access to emergency services.

Although Blount announced this initiative in April, the contents await final touches to a review of customer service guarantees for activities like connection, repairs and billing. The promise is that the Customer Charter will arrive soon, alongside the phone bill. Believe it when you see it.

Liz Fell

Austel

Robin Davey Steps Down

As Austel Chairman, Robin Davey, prepares to step down at the end of this month, the two leading candidates to succeed him face some daunting tasks. The toughest assignment will be to set new long-term objectives to facilitate the greater levels of competition and media convergence expected in the Australian telecommunications marketplace beyond 1997. But also looming is the possibility that Austel may be scaled down, or swallowed by a mega-communications regulatory body.

The two names most often mentioned in industry circles as possible successors are Michael Hutchinson, a long-standing career bureaucrat and Deputy Secretary, Communications, at the Department of Communications and the Arts, and Austel Board Member, Neil Tuckwell, whose background includes senior positions with OTC and Consolidated Press and a founding role with Clear Communications in New Zealand.

Whoever wins — the decision has apparently been made and an announcement from Communications Minister, Michael Lee, is imminent — will have rather large shoes to fill in Davey's absence. A former public sector lawyer and administrator, Davey became founding Austel Chairman five years ago and has seen at least as many Ministers come and go in that time.

As Davey sees it, the essential task for Austel is to minimise regulatory dependence among the carriers, value added service providers and future resellers. But this is like walking a tightrope, he says. "Move too far in the direction of commercial outcomes and you may end up in the less-than-satisfactory position of having the courts decide what are quite complex technical and economic issues. Move too far in the direction of the regulatory intervention and you run the risk of regulatory dependence and regulatory intervention becoming a substitute for real competition."

Asked about the potential impact of the Hilmer Report on Austel, Davey said he agreed with some elements, but had difficulty accepting others, such as some of the merger proposals. "While issues like price capping and price control might fit neatly into this scenario, there are other things like numbering plans, international and domestic standards and the Land Access Code that remain squarely within Austel's brief. The answer may be driven by convergence itself. There may be synergies in looking at the regulation of 'communications' as a whole, rather than the separate strands."

According to ATUG Executive Director, Wally Rothwell, "Austel has shown considerable initiative and backbone. It tackled the dominance issue and the COT cases without fear or favour. The carriers will always have arguments with Austel, but Austel isn't there to be our friend. It's there to facilitate deregulation and they're doing it pretty well." His major criticism was that while Austel had sometimes been slow to come to conclusions, at other times it had given ATUG far too little time to prepare its response to important issues. "But what it has done is make sure the community was well and truly consulted and that's very beneficial," he said.

Robin Davey

According to Coalition Communications Spokesman, Senator Richard Alston, the Australian marketplace will need Austel at least until 1997 and very probably until the year 2000. "Rather than being subsumed by some macro organisation like the US Federal Communications Commission (FCC) or the proposed Australian Competition Commission, where you'd have to build up niche regulatory groups in particular areas, my feeling is that we might as well retain the direct access that we now have through organisations like Austel." Praising Austel for doing "a pretty good job," Alston said the regulator had been hamstrung by Government inactivity and by being understaffed.

As for the future, Robin Davey would only say he had "a few irons in the fire" and was considering offers from the public and private sectors. While some industry figures claimed it has been an open secret for months that Davey would not be reappointed, the outgoing Austel Chairman replied: "When I took the job, I said I'd give it five years. The real risk if you overstay your turn is that you lose your objectivity. Better to quit while you're ahead."

Bernard Levy

We deliver CCITT, CCIR, ETSI, IEC EIA/TIA, JEDEC and IEEE

We can supply essential technical standards from these and other international organisations.

We can also supply a range of overseas newsletters and market and technology studies for telecommunications professionals covering:

- International Services Digital Network
- Synchronous Digital Hierarchy
- Asynchronous Transfer Mode
- Virtual Private Networks
- International Resale
- Digital Video Compression
- Virtual Reality
- Fibre Distributed Data Interface
- Fibre to the Home

We offer a personal service backed up by industry and technical expertise.

Call Sally Lloyd
on (02) 555 7377.
Fax (02) 818 2294.

STUART CORNER
INFORMATION SERVICES

PUBLICATIONS FOR THE AUSTRALIAN TELECOMMUNICATIONS INDUSTRY · NEWSLETTERS · RESEARCH REPORTS · INTERNATIONAL STANDARDS

FLEX

Motorola's FLEX™ will soon take the paging industry to new heights.

FLEX

The technology behind FLEX promises to revolutionise the paging industry, and take it to new heights. It offers more speed, more capacity and more features to transform paging into the communication tool of the future. FLEX will give the world smaller pagers (with longer battery life), pagers that communicate more voluminous data and pagers with better signal sensitivity. It will also provide greater opportunity for paging operators. FLEX, the new-generation high speed paging protocol: another breakthrough by Motorola, the world leader in telecommunications technology.

MOTOROLA

There's a *stack* of good reasons to *choose* HP Hubs.

EXPANSION SLOT
To add SNMP Management and security.
FUTURE EXPANSION
Memory, WAN bridge/routing and future high speed technologies.
DISTRIBUTED MANAGEMENT
Manage 16 hubs or 784 nodes up to 185 metres apart from a single managed hub.
EMBEDDED SAMPLING AGENT
For advanced traffic analysis.
FREE SOFTWARE
Free Windows based HP Stack Manager software included with every hub.
LIFETIME WARRANTY
Including first five years on site.
HEWLETT PACKARD

Industry compatible

Affordable

yes
NetWare
Tested and
Approved

HP, inventor of the stackable hub, now brings you the next generation. As you can see, there are now even more reasons to make sure you specify the market leader – HP.

The innovative, cost saving, new HP AdvanceStack hub is also standards based, and carries a lifetime warranty with a 5 year HP on-site warranty.

SPECIFICATIONS

- *Available in 12, 24 or 48 ports*
- *Manage up to 784 nodes with a single managed hub up to 185 metres apart*
- *Eavesdrop protection and access prevention security can be added with a management card*
- *Supports a range of media types including UTP, co-ax and fibre-optic with internal tranceivers*
- *Add bridge/routing and high speed technologies via the expansion slot*

For more information, see your nearest HP dealer or phone 13 1347. Or, to receive a fact sheet, phone our HP FIRST fax service on (03) 272 2627 and request ID number 7395.

HEWLETT® PACKARD
A Better Way.

JWT HPC 387

SNMP — DMI's Acid Test

DMI developers are judging the merits of two schemes that let the desktop specification share data with SNMP consoles.

Network managers won't rest easy until they can really live up to their titles — and that means gaining centralised control over all the gear on their nets. True, they can oversee bridges, routers, and hubs from an SNMP console, but for the most part PCs and the applications that run on them are out in the cold. PCs have limited resources to devote to management, and their operating systems aren't built to handle protocols like SNMP. Thus, even though desktop management systems are available, they remain proprietary.

But if leading vendors have their way, the desktop may soon be just as manageable as anything else on the network. Hardware and software that comply with the Desktop Management Interface (DMI), a set of APIs (application program interfaces) giving PCs a standard means of reporting information about themselves to management applications, are expected late this year.

If these products work as planned, DMI could open the way to managing PCs with the same tools that now monitor and control the rest of the enterprise — including SNMP management platforms like HP's OpenView, IBM's NetView/6000, and Sunconnect's Sunnet Manager.

For all its promise, though, DMI has a long way to go. DMI-compliant prototypes have been demonstrated for a while, and product rollouts should help convince vendors to add the spec to PC components like modems and fax cards and proprietary desktop management systems. But even though many suppliers support the spec publicly, privately they express reservations.

By most accounts, the acid test for DMI is going to be SNMP.

Right now, DMI is strictly a local interface, designed to gather data from a PC's operating system, applications, CPU, network adaptors, expansion cards, and attached peripherals. It furnishes a mechanism for storing this data and offers generic links to management applications. But it doesn't specify exactly how this data is to be delivered to a remote management console. In fact, DMI developers disagree as to the best way to link the spec with SNMP. Some argue that DMI and SNMP should share a common agent at the desktop. Others put the management burden at the server, with a gateway between the PC and the remote console.

DMI Details

Introduced nearly two years ago, DMI is the handiwork of the Desktop Management Task Force (DMTF), a vendor consortium that includes DEC, Hewlett-Packard, IBM, Intel, Microsoft, Novell, Sunconnect, and SynOptics Communications.

The main element of DMI is its agent software, also known as the service layer interface. This is a program that gathers data about PC hardware and software from ASCII text files — known as Management Information Files (MIFs) — embedded in the PC operating system, peripherals, and network adaptors. The agent can run as a standalone application under multitasking operating systems like OS/2 or as a TSR (terminate-and-stay-resident) program under DOS or a DLL (Dynamic Link Library) under Windows.

MIFs are similar to the MIBs (management information bases) used with SNMP, although far simpler (MIBs are written in ASN.1 — Abstract Syntax Notation.1 — an arcane data definition language). The MIF for a modem card, for example, might contain data about the standards it supports.

A MIF-to-MIB translation program of some sort will be required in most products that support both DMI and SNMP. In some instances, these programs will convert every item in a MIF file into a MIB format. Thus, every element of the PC would have its own separate MIF and MIB. In other cases, a utility will convert information from multiple MIFs to information resident in a single MIB, sometimes nicknamed a 'MIF MIB.'

Some DMTF members have been experimenting with various MIF-to-MIB conversion techniques. Novell may provide a MIF MIB for the clients that work with its NetWare Management System (NMS). In

DMI Plus SNMP: Call My Agent

The big question for DMI is how it will be integrated with SNMP. Advocates of the agent approach want to put DMI and SNMP agents in the same PC, with all DMI data being converted to SNMP before it's sent to a central console (A). The server approach puts DMI at the desktop and an SNMP proxy agent at the server to convert data and pass it to the console using TCP/IP, IPX, or some other LAN protocol (B).

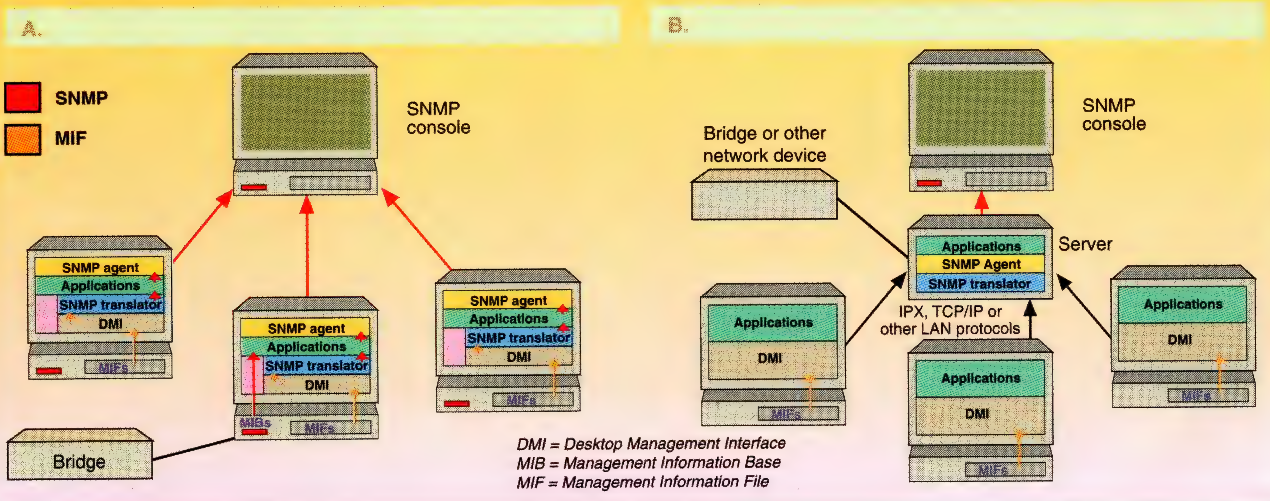

DMI = Desktop Management Interface
MIB = Management Information Base
MIF = Management Information File

Market Watch

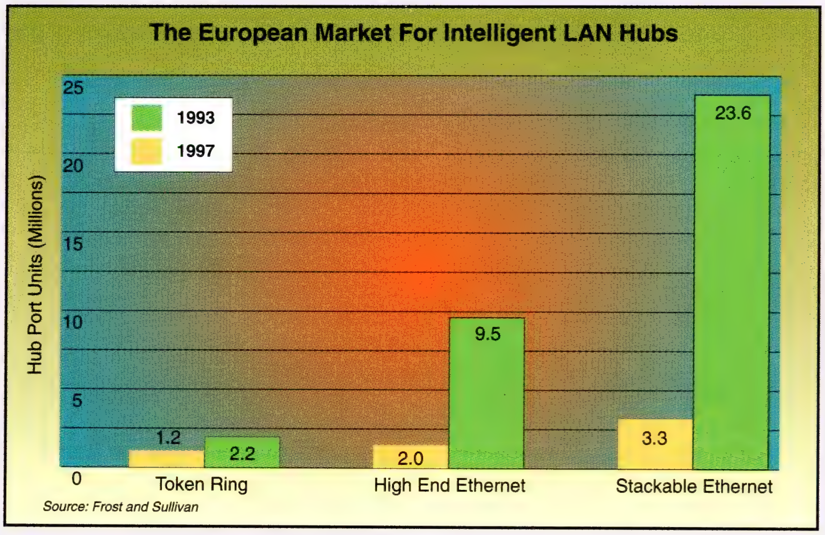

The European Market For Intelligent LAN Hubs

Source: Frost and Sullivan

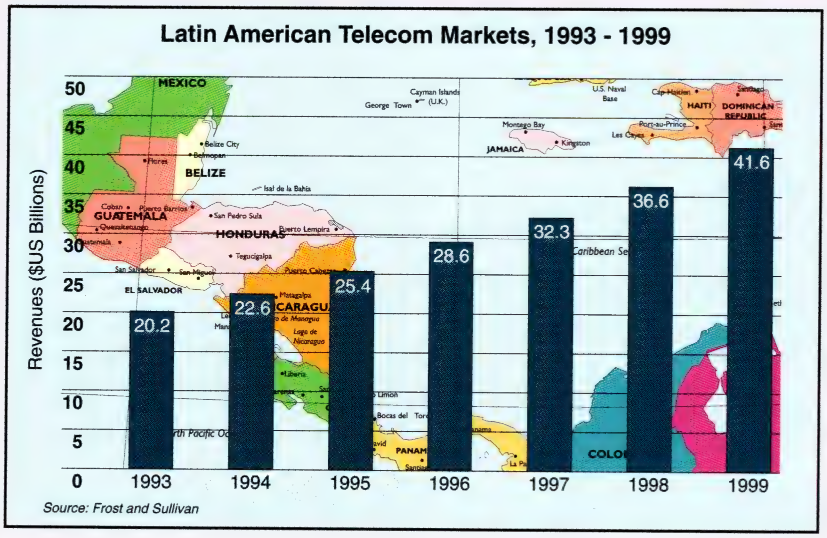

Latin American Telecom Markets, 1993 - 1999

Source: Frost and Sullivan

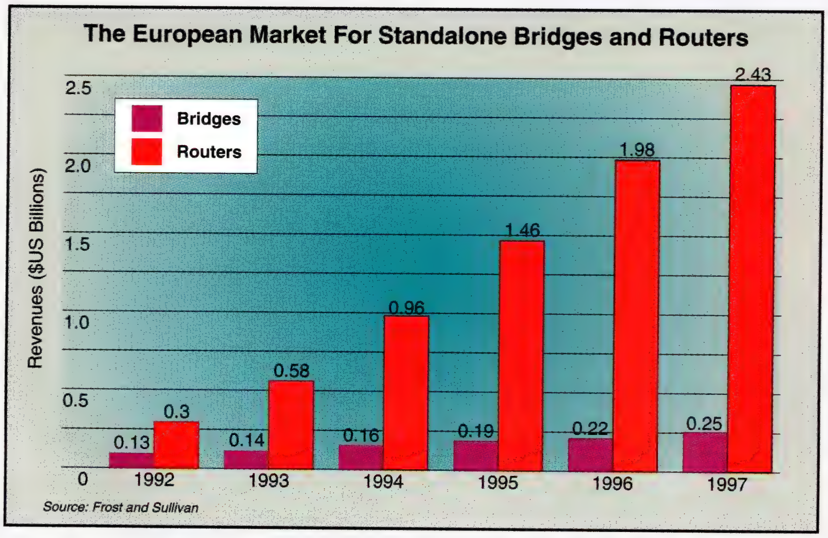

The European Market For Standalone Bridges and Routers

Source: Frost and Sullivan

the future, these clients will gather data from PCs via DMI and convert it to SNMP, which will then be picked up by a remote NMS console.

The DMI agent code, specifications, and related documents for Windows, Windows NT, OS/2, and DOS are available from the DMTF (Tel: +1 (503) 221 2945). Users also can obtain DMI code and specs from Microsoft's bulletin board or from the Internet (via anonymous FTP aurora.intel.com in the pub/DMTF directory).

There's just one catch — albeit a big one. Vendors must bundle MIFs with their products if DMI is going to work. Microsoft — a DMTF mainstay — says it will, but others may need to be persuaded.

Agents at Work

The DMI agent accesses a MIF through a set of APIs called the component interface. And it stores MIF information in a database on the PC. Desktop management applications use DMI APIs to query the agent, which responds by opening up the database and releasing the requested data. For instance, an inventory package may want data about the type of network connections and peripherals supported by a particular PC. The DMI agent would check its MIF database, find the information associated with network adaptor, fax, and modem cards and pass it along.

At this point, though, DMI offers no transport protocol to transmit data from a PC to a remote management station. In fact, DMI doesn't support any specific management protocols at all — a benefit, the DMTF says, because it allows the spec to be implemented with SNMP, CMIP, and proprietary desktop management packages.

But this sort of openness comes at a price. The only way for DMI to be integrated into desktop management software is with vendor APIs. In other words, suppliers must write DMI into their products. Most suppliers have proprietary desktop agents integrated with operating systems like Windows, NetWare, DOS, and OS/2. These agents gather data on PCs as they boot up or log on to a server. The DMI agent won't necessarily replace them or improve on the quality of information they already provide. But it will offer a way to store inventory data in a common database on a PC that can be accessed by different vendors' desktop management applications — something that's virtually impossible today.

Most vendors in this market say that adding a DMI agent to their wares won't be a major undertaking. But many companies are taking a wait-and-see attitude toward.

Dovetailing DMI

The real challenge, though, will be tailoring DMI to work with SNMP. DMTF members say that a lack of specific support for SNMP doesn't mean that DMI is unable to work

The MIF Match Up

The MIFs (Management Information Files) used by DMI deliver more detailed information about a PC's inventory of hardware and software than can be obtained via SNMP. This is because MIFs will be written by the manufacturers of specific products. For example, a vendor of plug-in cards could include warranty date and the date the card was installed in the PC as part of its MIF. Further, an adaptor card MIF might offer data from special diagnostic registers known only to the manufacturer.

In contrast, systems based on SNMP typically use the host resources MIB (management information base) and other general-purpose mechanisms to gather data from a computer. This MIB relies on configuration data that's easily accessible by scanning a PC when it starts up — CPU type, memory allocation, network adaptor cards, and so forth.

DMTF members think it's vital that vendors bundle MIFs — the ASCII files that describe a particular desktop component to the DMI agent — into their PC products. To promote the adoption of MIFs, the DMTF has formed working groups in which leading vendors of PC hardware and software define the formats and specify what information MIFs should contain. Working groups have outlined MIF parameters for cards and printers, with committees for fax cards and modems still being organised.

Another group comprising PC vendors is writing a MIF that will allow DMI agents to identify characteristics of disk drives, motherboards, memory, buses, and other features. (Vendors like Microsoft will work on their own to provide MIFs for desktop operating systems.)

A working group also has been proposed to define MIFs for devices that use the as-yet-unreleased Plug 'n Play interface created by Microsoft and Intel. This spec allows an adaptor card to install itself on a PC automatically by interacting with operating system drivers. Microsoft has said that Plug 'n Play will make its debut in the Chicago operating system, the follow-on to Windows 3.1 that is slated to ship during the second half of the year.

There's been an effort on the part of the DMTF to match its work with other industry groups. The printer MIF working group, for example, has basically the same membership as the printer MIB group at the IETF. MIF groups also use the IETF host resources MIB.

Another handy reference is the Requirements Specification document of the Distributed Support Information Standards Group (DSIS), a committee of users and vendors looking to define all the computer hardware and software components that need to be managed in a distributed network.

Mary Jander

with it. In fact, many DMTF members and affiliates have demonstrated SNMP-plus-DMI prototype products which use DMI to collect data from PCs and rely on SNMP to transport it to the central console for display. Most of these products, though, have yet to hit the streets.

Every vendor trying to meld DMI and SNMP has its own take on how this is to be done, but generally speaking the schemes fall into two categories. The so-called agent approach calls for DMI and SNMP agents to coexist in the PC (see the figure on page 31). The server approach calls for only DMI at the desktop, with an SNMP proxy agent converting data supplied by the DMI agents into a format that can be read by the management station. Even though an SNMP console would be employed, a range of transport protocols could be used to send management information — including TCP/IP or IPX.

Advocates of the agent approach include IBM, Novell, and SNMP Research. These vendors say they'll add DMI to SNMP agents. In most cases these agents will translate DMI MIFs into SNMP. Proprietary APIs would be used to allow desktop management systems to access information contained in the agents.

Another agent advocate, Netmanage, is taking a different approach. The vendor plans on using the Windows SNMP APIs by the Windows SNMP Task Force, an ad hoc industry group whose members also participate in the IETF (Internet Engineering Task Force) and the DMTF. (These specs are available on the Internet via anonymous FTP to the WinSNMP directory of sunsite.unc.com.) Currently, the WinSNMP APIs, initially developed by Netmanage and Microsoft, support Windows management consoles; Netmanage has submitted a set of agent APIs to the Windows SNMP Task Force.

Those who favour the agent approach point to the fact that the DMI specs currently lack any specific protocol for remote management. Putting SNMP on each PC would allow it to be managed just like a bridge, router, or hub. Further, equipping a PC with its own SNMP stack allows it to communicate directly with a central console — there's no need for a server or hub to act as intermediary. This approach is thus not tied to a specific type of server, such as one from Microsoft or Novell.

SNMP at the desktop also gives net managers access to performance information for diagnostic purposes. For example, a manager could use an SNMP agent on a PC to filter network packets and then compare the performance statistics with information collected by an agent at a hub.

In fact, a PC application that understood both SNMP and DMI could be set up to act as a mid-level manager, filtering inputs from other PCs on a segment as well as from local devices and selectively passing along data to a central console. SNMP Research has demonstrated workstation-based mid-level managers that do just this. And this scheme is more economical than using a hub to do the same job, since a hub hardware module is needed to support the SNMP agent. Companies like Netmanage bundle free SNMP agents with their TCP/IP protocol stacks.

The Server Side

Those who favour the server approach say it will give users the best of DMI and SNMP without having to support them both at the desktop. They say SNMP would still be used to manage servers, which also could selectively screen DMI data from attached workstations, convert it to a special SNMP MIB format using an SNMP proxy agent or proprietary APIs, and then send it via a gateway or similar mechanism to an SNMP console. This approach would also eliminate polling between an SNMP management station and remote PCs.

The server approach seems to be favoured by Microsoft. While that vendor hasn't publicly stated its plans for the DMI, it's indicated that DMI will be part of its new desktop management system, code-named Hermes, which is slated for release soon.

Microsoft's strategy seems to be as follows: DMI agents integrated into DOS, Windows, Windows NT, and OS/2 would send information to Hermes at a server equipped with Microsoft's Windows NT Advanced Server software. Once it arrived, the DMI data would be converted to SNMP for use by remote management stations. It would also be placed into the server's Microsoft SQL Server database, where it would be accessible to applications from Microsoft and third parties. It's important to note that Microsoft's DMI strategy won't be fully implemented in the first release of Hermes, which will only contain local CMI agents. SNMP compatibility will be added later, Microsoft sources say.

Intel says that it will add a DMI agent to its LANdesk management suite, which runs standalone or under Novell's NMS or HP's OpenView. A LANdesk-to-SNMP gateway for Novell servers will arrive this year.

Server proponents say that SNMP may not always be the favoured technique for transporting data from a desktop to a console; RPCs (remote procedure calls) could prove more efficient in the long run. It's easier to convert DMI data at a server to multiple formats than it is to revise multiple PC agents. They also claim that combining SNMP and DMI in the PC overloads it.

Some observers seem disappointed with both the agent and server techniques because they both require DMI data to be converted to SNMP. These folks would like to see the DMI's management interfaces adapted directly for use with SNMP management stations. DMTF sources say that within the next couple of years, a new set of DMI specs with better SNMP support might surface. First, though, the DMTF wants to promote the DMI without favouring one particular management scheme. Once the DMI is accepted, DTMF members claim, there will be time to address a closer union with SNMP.

Mary Jander

Network Security

What is Your LAN Vendor Doing About Security?

This past February, while most of Norway was preoccupied with getting ready to host the Winter Olympics, intruders broke into the National Museum of Art in Oslo and made off with one of the country's most highly-prized art treasures: Edvard Munch's Expressionist masterpiece, 'The Scream.' In place of the canvas, the thieves left a note thanking the authorities for the very poor security that made their crime so easy to commit.

At about the same time the Munch heist was being pulled off, the US-based Computer Emergency Response Team (CERT) issued a warning that unauthorised users had broken into tens of thousands of Internet-attached networks. Once network crashers gained initial entry into networks — a little-known feature on some Sun and Solbourne operating systems served as a conveniently unlocked back door — they were able to capture user logins and passwords, ensuring easy and undetected access in the future. The end result, according to CERT: all network-connected sites were at risk from invasion.

The Internet interlopers may not have left any thank-you notes behind, but the message implied in their success rings just as clear: the security mechanisms now in place in most corporate nets are woefully inadequate, leaving sensitive data and applications wide open to all sorts or criminal mischief.

Security risks are especially troublesome for enterprise networks that encompass scores of geographically dispersed, internetworked LANs. The problem isn't just that sensitive data may be stored on any number of widely scattered servers; it's also that the operating systems driving those LANs initially weren't intended to handle the high-profile, enterprise-wide applications that they now run.

Fortunately, some much-needed help in curbing LAN intrusion is on the way. The four leading vendors of LAN operating systems — Banyan, IBM, Microsoft, and Novell — have started adding some sophisticated security features to their respective products. Previously, these features were available primarily as add-ons from third-party providers.

The security measures taken or planned by LAN operating systems vendors fall into four basic categories: the type of encryption scheme used for passwords, support for an extended password encryption scheme known as packet signature, the ability to manage security features from a central location, and expected support for Kerberos, a state-of-the-art authentication system developed by the Massachusetts Institute of Technology.

In terms of these four categories, each operating system vendor is taking a unique approach to security. For instance, although all four already have added some form of encryption to their operating systems, only two — IBM and Microsoft — have opted to install the Data Encryption Standard (DES) algorithm. And only Novell has decided to offer the packet signature encryption extension.

Make or Break

By all indications, suppliers of LAN operating systems are treating security as a make-or-break component in their efforts to sell their products as enterprise-wide networking platforms. Banyan, Microsoft, and Novell are seeking certification under the US National Security Agency's Controlled Access Protection C2 (Class 2) program.

These vendors are hoping that by going through the painstaking approval process, they can further convince users that they are doing all they can to keep networks intruder-free.

The rise in distributed computing has network managers more worried than ever about the safety of their systems. More than 70% of the 870 US network administrators responding to a 1993 survey by the consultancy division of accounting firm Ernst & Young said they believed that the risk of unauthorised incursions into their data nets had grown in the previous five years. More than half of the respondents said that in the two-year period covering 1991 and 1992, their companies lost money because of misdeeds by employees, while 43% said their organisations suffered financial losses as the result of meddling by outside hackers. About 10% of all those losses involved $100,000 or more, according to the survey respondents.

The US Federal Bureau of Investigation says network managers are right to think that computer crime is on the increase. "In 1990, we had 49 computer crime investigations pending," says Hal Hendershot, special agent with the FBI's economic crime unit. "Now there are 170."

Who Goes There?

The latest versions of the leading LAN operating systems — Banyan Vines 5.536, IBM LAN Server 3.0, Microsoft Windows NT and Windows NT Advanced Server 3.1, and Novell NetWare 4.X — all come with password management tools that let network administrators define the rules governing how users pick their passwords. With password management, administrators can specify that passwords be a minimum length, or that they be changed at regular intervals. (NetWare 4.X includes a feature that prevents users from recycling old passwords.) Managers also have the option of allocating passwords to users for greater control.

The latest operating systems also encrypt the user's password during log-on, to prevent eavesdroppers equipped with LAN analysers from capturing and copying passwords. Banyan's password encryption is based on proprietary technology, while IBM and Microsoft encrypt their passwords using the widely deployed DES algorithm.

Continued on page 38

LAN Security From the Big Four

VENDOR	OPERATING SYSTEM	PASSWORD ENCRYPTION	PACKET SIGNATURE	CENTRALISED MANAGEMENT	KERBEROS SUPPORT	C2 COMPLIANCE
Banyan Systems	Vines 5.536	Proprietary	No	Yes	No	Server compliance under evaluation
IBM	LAN Server 3.0	DES	No	For user definitions, not access rights	Planned for 1994	No
Microsoft	Windows NT and Windows NT Advanced	DES	No	Yes	Planned for 1995	Server and client compliance under evaluation as separate entities
Novell	NetWare 4.x	RSA	Yes	Yes	No	Server and client compliance under evaluation as a network entity

Jtec

Solutions

What Keeps BP Australia On The Move?

Jtec's communications network is BP's efficient business engine.

BP Australia's commitment to the highest standards of quality in its products, service and distribution, keeps the company on the move throughout the 90's and is earning it an enviable reputation.

Behind much of the company's success lies its sophisticated communications network, an intricate highway that links its operations across the country and around the globe. In this area, as with its products and services, BP Australia demand the best.

When it came to choosing a communications partner with an understanding of its business issues, the choice was simple—

Jtec.

BP Australia relies on its communications network. It is the key to maintaining an efficient business network. The key to maintaining BP's winning edge. That's why BP relies on Jtec.

Jtec. *Keeping BP Australia On The Move.*

Head Office
Jtec Pty Limited
ACN 003 169 088
Unit 3
118 - 122 Bowden St
Meadowbank 2114
NSW Australia

Phone

Sydney	61-2-809-6933
Melbourne	3-263-0200
Brisbane	7-229-9233
Canberra	6-253-3399
U.K.	44-494-473757
Europe	353-21-277422

Fax

Sydney	61-2-809-6619
Melbourne	3-263-0215
Brisbane	7-363-0023
Canberra	6-201-3003
U.K.	44-494-536254
Europe	353-21-270420

Full-Duplex Ethernet

More Niche Than Necessity?

The Ethernet switches now on the street can torque up the technology's throughput to 10Mbps per port, which should go a long way to taking some of the overwork out of today's networks. But this type of gear could prove to be only half of the Ethernet bandwidth story. A new generation of so-called full-duplex switches and adaptors can actually double that performance, maxing out at 20Mbps on a dedicated LAN segment.

What's more, there's no price penalty for full-duplex technology — but in order for the scheme to work, both a LAN switch and an adaptor are required. Many of the switch vendors that offer it do so as part of their standard feature set. And suppliers of full-duplex adaptors say their cards are priced to complete aggressively with their half-duplex counterparts.

Double the bandwidth without paying twice the price? If that sounds as if customers should be lined up 10 deep, think again. Full duplex's detractors point out that only multithreaded operating systems like OS/2 and Windows NT can take full advantage of this hardware. And they also comment that little has been done in the way of standardisation, which could lead to interoperability problems.

Right now, it looks as if most users will opt instead for the 100Mbps shared bandwidth of FDDI or simply wait for ATM. A recent survey of 500 network administrators by International Data Corp. (IDC) showed that only 4.8% of the respondents favour full-duplex Ethernet, compared with 36.9% for ATM and 35.3% for FDDI.

Niche or not, though, a growing number of vendors are delivering — or planning to deliver — full-duplex switches and adaptors. Almost all of the gear available thus far is for Ethernet, but the technology also is making inroads into Token Ring, FDDI, and Fast Ethernet. And many of these vendors are working on standards for full-duplex technology, although some in the Ethernet camp insist that concerns with interoperability are more the result of bad press than hard facts.

Two-Way Traffic

The way Ethernet is set up, only one node on any network segment can transmit data at any time. If two nodes try to send data simultaneously, the collision detection part of Ethernet's CSMA/CD protocol automatically tells one of the stations to back off and try again. This network control mechanism is the reason that bandwidth is restricted to 10Mbps on switched Ethernet connections: attempts by the end node to send and receive data at the same time are foiled since they are interpreted as collisions. But because full-duplex can only work when there is one end node per segment — an important restriction to understand — collision detection is unnecessary. And once this mechanism is disabled, nodes are free to transmit and receive data at the same time — bringing aggregate throughput to 20Mbps.

Full-duplex vendors see two major applications for their equipment — switch-to-switch connections and switch-to-server connections. In both cases, the idea is to use full-duplex to break the bandwidth bottleneck that can occur when a number of nodes try to send data simultaneously over the same switched connection.

But not everyone is a believer. Some critics doubt that the twofold increase in their performance will be adequate. Their arithmetic is simple: "If you have 10 10-Mbps Ethernet connections feeding into one 20Mbps full-duplex Ethernet connection, you're still going to have a problem," says Dave Husak, systems architect at 3Com's switching division (formerly Synernetics). The approach favoured by 3Com and other switch vendors is to use a higher-speed technology like FDDI for switch and server links and switched Ethernet between network nodes.

Loose Threads

Critics of full-duplex are quick to point out that for nodes to take full advantage of the technology, they must be running a multithreaded operating system. In essence, these OSs divide tasks into multiple threads (streams of instructions) that run concurrently. The key, according to Craig Johnson, Director of Marketing at Lannet in the US, is that the OS must be "capable of handling traffic streams in and out at the same time." OS/2, Windows NT, and NetWare 3.11 and 4.X (server versions) are all multithreaded operating systems, as are different varieties of Unix, including SCO Unix and IBM AIX, Nodes running under

Putting LAN Data on a Two Way Street

Full-duplex Ethernet can double the throughput of switch-to-switch and switch-to-server connections. The technology disables the collision detection portion of the CSMA/CD protocol, allowing for two way traffic.

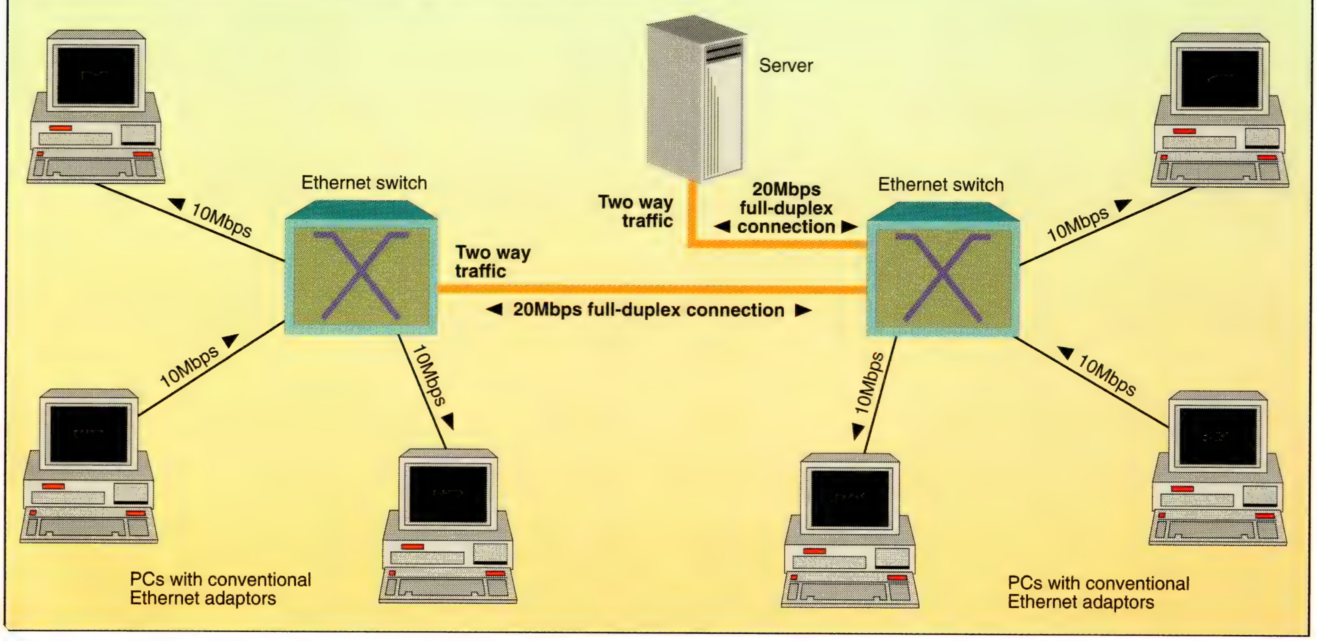
Server
Ethernet switch
Two way traffic
20Mbps full-duplex ◄ connection ►
Ethernet switch
◄ 10Mbps
10Mbps ►
Two way traffic
◄ 20Mbps full-duplex connection ►
10Mbps ►
10Mbps ▼
10Mbps ▼
◄ 10Mbps
PCs with conventional Ethernet adaptors
PCs with conventional Ethernet adaptors

DOS and Windows will at best see a marginal performance increase from full-duplex hardware.

Of course, full-duplex Ethernet has one big advantage over other bandwidth-boosting schemes: it's a lot cheaper. Of the four vendors now shipping full-duplex Ethernet switches, Cabletron Systems, Nbase Switch Communications, and Netwiz, say that full-duplex is standard on their boxes. The Etherswitch from Kalpana costs extra.

Vendors of full-duplex Ethernet adaptors like Cogent Data Technologies, Compaq, Digital Equipment Corp., and IBM indicate their products are priced to sell against half-duplex cards.

A small minority of users will even be able to upgrade installed adaptors rather than swapping in new ones. It all comes down to whether the Ethernet controller on the adaptor card features dual FIFO (first-in, first-out) paths that can send and receive traffic at the same time.

If that's the case, upgrades can be accomplished via software downloads, which can translate into very big savings. First off, no new adaptors need be purchased. Then, there's no need to open up the PC and install a new adaptor, a task that can take anywhere from a few minutes to several hours.

Cards carrying the latest 82596 Ethernet controller from Intel or the Super Eagle Ethernet controller from Texas Instruments should be upgradable. Unfortunately, none of the Ethernet controllers from Advanced Micro Devices (AMD) or National Semiconductor can be upgraded.

National Semi's silicon represents by far the largest percentage of the controllers now in place on Ethernet adapters. Both manufacturers have full-duplex chips on the drawing board.

Point, Counterpoint

Those critical of full-duplex point out that the lack of standards could lead to interoperability problems. Its backers, however, are quick to deny the charge. They claim that since full-duplex mainly involves disabling collision detection — rather than adding something new — it's easy to get products to work together. "We have tested our switch without any problem with cards from Compaq, DEC, and Seeq," says Yechiel Kurtz, Design Engineer at (Israel-based) Netwiz.

Further, about a dozen vendors banded together last year to form an industry group to address interoperability and other issues. The Full Duplex Consortium is only involved with the further development of full-duplex Ethernet. It's also working to define an auto-detection (or auto-negotiation) option that will allow full-duplex Ethernet adaptors and switches to sense the presence of a full-duplex device at the other end of a connection and automatically move into full-duplex mode.

Technology Update

■ AT&T, Novell Announce Global Business Network

AT&T and Novell have announced an alliance to provide a new service, to be known as AT&T NetWare Connect Services. The new network will support both Novell's IPX and the TCP/IP protocols and will allow users to access a range of AT&T's public data services such as AT&T Network Notes (a Lotus Notes application) and AT&T Easylink messaging service, as well as applications available on the Internet. This announcement comes shortly after Novell and AT&T's joint announcement of NetWare Telephony Services, a software platform that provides the ability to link existing telephone and computer networks, and in conjunction with PABX systems can deliver sophisticated call control and voice processing capabilities.

■ IETF to Introduce New E-Mail Draft Standard

The Internet Engineering Task Force is set to produce its first standard for accessing and manipulating electronic mail. The organisation is nearing completion of a draft specifying an Internet Message Access Protocol (IMAP), which covers commands for retrieving messages and manipulating mailbox activity in TCP/IP networks. IMAP will provide standard ways of creating, deleting, renaming and checking the status of individual mailboxes, which are to be organised into hierarchies of mailbox folders on e-mail servers in TCP/IP networks. The draft will also provide a way for TCP/IP networks to handle e-mail from a variety of network types.

■ Hughes 'Revolutionises' Domestic Phone Services

Hughes Network Systems has launched its Cable Telephony Transport System (CTTS), which the company claims will revolutionise the delivery of domestic telephone services. The system coverts telephony signals into a radio frequency (RF) carrier for distribution through existing fibre and coaxial TV cables. It will allow cable service operators to provide a competitive phone service and to offer fax and business data services at up to 2Mbps. The system is comprised of a 'head end' unit that modulates signals for digital transmission over the RF carrier to kerbside or switching terminals, and a remote user interface (RUI) which demodulates the signal for delivery via conventional twisted pair. Officials said each kerbside RUI can support up to 30 subscribers up to one kilometre distant.

■ Simplifying Secure SNMP

Developers working on SNMP 2 are hopeful that a new set of algorithms will relieve some of the hard-slog involved in deploying Secure SNMP. The new algorithms are designed to simplify the configuration routines required to monitor and control network devices. Currently, configuration files have to be set up for each SNMP console and agent on the network, and each time a device is added or removed all Secure SNMP configuration files have to be updated. The new algorithms will greatly simplify the update process by ensuring a new device's configuration file links a network manager's user ID and password with the Secure SNMP configuration coding used to protect the device from unauthorised tampering.

■ Document Management Systems

An industry group called the Enterprise Document Management Coalition has developed a set of APIs designed to streamline the management of documents created in different formats. The new APIs are known collectively as Shamrock, and provide a format in which documents from a wide range of applications can be stored for access via a range of e-mail or file transfer techniques. Meanwhile, Novell and Xerox are collaborating on the development and marketing of Document Enabled Networking, which is also designed to aid document management over enterprise networks, and can locate documents regardless of storage location or format.

■ Stanilite, InterDigital Sign CDMA MoU

Sydney-based Stanilite Electronics and InterDigital Communications Corporation (IDC) of Pennsylvania have signed a Memorandum of Understanding covering the joint marketing and manufacture of IDC's Broadband Code Division Multiple Access (B-CDMA) wireless technology. Under the terms of the agreement, InterDigital will provide B-CDMA engines, technology and mobile handsets, while Stanilite will provide network infrastructure including microwave links, switches, base stations and fixed subscriber terminals.

■ Novell Sets Up Anti-Piracy Hotline

Novell Australasia has established an Australia-wide, toll-free hotline in a bid to combat software piracy and trademark and trade name infringements by resellers and end users. The number for the hotline is 1 800 627 554. Callers outside Australia can ring +61 2 925 3131.

Until some form of auto-detection is specified, however, users will have to manually set jumpers on adaptor cards to enable full-duplex communications. Once again, this is a tedious, time-consuming task. "We are optimistic that the consortium should finish its work by February," says David Solomon, product marketing manager of the Ethernet division at National Semi.

Moves are afoot to have the consortium's work on auto-detection included within IEEE 802.3. "There have been informal discussions," says Michael O'Connor, Senior Engineering Manager at DEC.

The IEEE is also scheduled to look at full-duplex Token Ring. Insiders expect the organisation to rubber-stamp a draft version of a full-duplex Token Ring spec for inclusion in IEEE 802.5. "The idea is to extend the MAC-level protocol so that adaptors can act in full-duplex or half-duplex modes," says Norm Strole, senior engineer at IBM's Networking Hardware Division, and one of the proposal's developers. He says full-duplex will be included as an option in 802.5 and will feature an auto-detection feature similar to the one under development for Ethernet.

No work has yet been undertaken to standardise use of full-duplex with either fast Ethernet or FDDI.

The Final Frontiers

While almost all of the full-duplex products now available are based on Ethernet, IBM is leading the charge into another market: full-duplex Token Ring. IBM also is one of four companies shipping a full-duplex Ethernet adaptor. Clearly, its unequivocal endorsement of full-duplex gives the technology's prospects a significant boost.

IBM is shipping a range of full-duplex Token Ring adaptors, as well as a software package that allows any of its 4/16Mbps Token Ring adaptors to be upgraded to full-duplex service. Unfortunately, users may have to wait until the third quarter before they can take advantage of the new capability; that's when Big Blue is expected to begin deliveries of an eight-port full-duplex Token Ring switch.

On the FDDI front, Cabletron Systems says it will launch both 12 and 24-port full-duplex FDDI switch modules for its MMAC and MMACplus hubs soon. The vendor says it also has plans for full-duplex FDDI adaptors.

SynOptics indicates that it will introduce a full-duplex fast Ethernet switch sometime this year. The device will support full-duplex fast Ethernet connections on some ports and conventional Ethernet connections on the others. Industry sources also suggest that Intel is planning a full-duplex fast Ethernet adaptor, while 3Com is rumoured to be working on a full-duplex fast Ethernet switch.

Stephen Saunders

Network Security from page 34

Novell encrypts passwords using the RSA algorithm developed by RSA Data Security.

DES is known as a symmetric encryption scheme. This means that the sender and receiver use the same key to encrypt and decipher messages. RSA is an asymmetric scheme, in that messages are sent using a public key and deciphered using private keys unique to each user.

Sign Here

Novell has taken encryption a step further by adding the packet signature facility to NetWare 4.X. Packet signature, developed by Novell in response to an attack by a group of hackers based in the Netherlands, attaches an encrypted code to every packet on the network. The code identifies which client/server session each packet belongs to; it allows the operating system to monitor traffic and ensure that sessions cannot be interrupted and taken over by an imposter.

The extra processing power and additional packet size required by packet signature degrades network performance by about 8%, according to Novell. The packet signature feature is optional; net managers choose whether to activate it at start-up.

Novell touts packet signature as an extra layer of security. But Microsoft asserts that Novell had to institute packet signature to plug a hole in NetWare that allows interlopers to hijack client/server sessions. Microsoft says the NetWare loophole is caused by Novell's use of its connectionless IPX network protocol.

In contrast, Microsoft's Windows NT uses TCP/IP to carry client/server traffic. Microsoft claims it is much more difficult to hack into a TCP/IP session because IP is a connection-oriented protocol. Each packet is transmitted in a numbered sequence, and acknowledgments are sent for each one. "We don't have the same problem [as NetWare] because IP is a secure protocol to begin with," says Ty Carlson, Product Manager for Windows NT TCP/IP technologies at Microsoft.

On the plus side for Novell, the vendor says it is planning to extend NetWare 4.X's RSA encryption to the data portion of every packet — rather than just for passwords and packet signatures. To achieve the same level of encryption security, net managers currently must buy third-party products — a move that can run into some serious spending.

Along with using encrypted passwords to ensure that users are who they say they are when they log on to the network, the latest LAN operating systems from Banyan, IBM, Microsoft, and Novell all support single sign-on — the ability to check users' authorisation every time they access a new network resource. The operating systems perform these additional checks in the background, meaning that users don't have to enter a new password every time they need access to a different source.

All four operating systems offer several other tools aimed at weeding out possible intruders. For instance, they all contain a feature that can be activated to prevent a single user from logging on to more than one workstation at a time. Access by individual users also can be restricted to workstations with specific addresses.

As a further measure, all four have a scheduling facility that lets net managers specify the times of day when users are allowed access to the network. IBM and Novell now include a feature that automatically logs users off the network if there is no activity from their keyboard for a specified period of time.

Enterprising Management

As suppliers of LAN operating systems outfit their offerings with improved security, they need to make sure that those features don't become too cumbersome to implement. With older LAN operating systems, such as NetWare 2.X and 3.X and Microsoft LAN Manager, network administrators have to assign resources to users individually. The only way to do that is to go to each server and modify the access information stored on each one. That's fine for single-site networks with a few servers and a few dozen users, but for enterprise LANs with scores of servers and hundreds or thousands of users, that approach just won't fly.

All four leading LAN operating system vendors have put substantial work into making security implementations easier for users. Most notably, the latest releases from Banyan, Microsoft, and Novell allow managers to view the network as a single entity rather than as a succession of servers. This single-view perspective enables administrators to allocate access rights centrally from a single console, without having to alter rights on each server individually.

Access rights can be allocated on a group basis. For example, a network manager can give all the users in a marketing group access to a list of the same resources. To give a new user access to the same rights, the manager simply adds that user's name to the group list. Taking away a user's access rights is just as easy: simply delete the name from the list.

IBM has some catching up to do in this area. LAN Server's Domain Controller facility allows net managers to define which users (and passwords) will be recognised by a network from a single console, and then distributes them across multiple servers. However, the net manager still has to define which files each user can access on a server-by-server basis.

"Our users don't often complain about not being able to centrally allocate access rights," says Bob Blakley, LAN Systems Security Architect at IBM's Personal Software Products LAN Systems Division. However, he admits that LAN Server's inability to deliver a list of which resources can be accessed by each user has drawn negative comments from IBM users.

Both IBM and Microsoft are planning to significantly enhance the security facilities of their respective network operating systems through support for Kerberos, an authentication system developed as part of MIT's Project Athena. Kerberos uses a dedicated, physically secure server to provide a two-stage, encrypted authentication procedure for every user and every resource on a network.

Microsoft says Kerberos will be a feature of Cairo — the object-oriented version of Windows NT scheduled for release in the first half of next year. Meanwhile, IBM says its Kerberos implementation will come as part of a new version of LAN Server. That new version will support the Distributed Computing Environment (DCE) from the Open Software Foundation (OSF). OSF has adopted Kerberos as one component of DCE. IBM expects its forthcoming DCE support to improve LAN Server's security abilities in two other ways. "The administrative model for LAN Server needs enhancements," says Trudy Henke, DCE Brand Manager at IBM's Personal Software Products Division. "We are changing it so that it is closer to the DCE model."

DCE features a standard security API developed by the Internet Engineering Task Force that will enable LAN Server to trade data about authentication and access rights with other software supporting the IETF API. Currently, all local area network operating systems carry security information in proprietary data formats, preventing them from exchanging information.

A version of LAN Server that supports DCE is expected to be in beta test by the end of the year, says IBM. Banyan and Novell say they have no plans to implement Kerberos. However, Open Computing Security Group Inc. a US-based security consultancy and vendor of network security software products, is planning to launch add-on products that support Kerberos on Vines and on NetWare later this year.

Test Time

Despite the addition of many new features to their LAN operating systems, when it comes to security many network managers would rather not take a vendor's word that its product offers watertight protection. For that reason, some vendors have entered their products for evaluation by the US National Security Agency (NSA).

The NSA offers stringent testing based on a program originally devised to evaluate the security features in computer systems protecting classified information for US Government departments, like the Department of Defense. The actual tests are conducted by the National Computer Security Centre (NCSC). NSA defines seven levels of security, which are laid out in a tome called the Trusted Computer System Evaluation Criteria (TCSEC) handbook — also known as the Orange Book. The seven levels range from D (the lowest level of security) to A1 (the highest level). Banyan, Microsoft, and Novell are seeking NSA's Controlled Access Protection C2 (Class 2) rating. C2 defines four key requirements:

- Discretionary Access Control — The operating system must provide a mechanism for authorising access between a user or group of users and 'named objects,' such as directories, files, and programs;
- Object Reuse — When a storage device such as a hard disk or tape drive is reassigned from one network user to another, the operating system must cleanse the medium to ensure that the new user cannot access any of the information stored there by the previous owner;
- Identification and Authentication — The operating system must require users to

SICK OF BEING BEATEN BY BUGS & VIRUSES? LET **INS** PROVIDE THE CURE

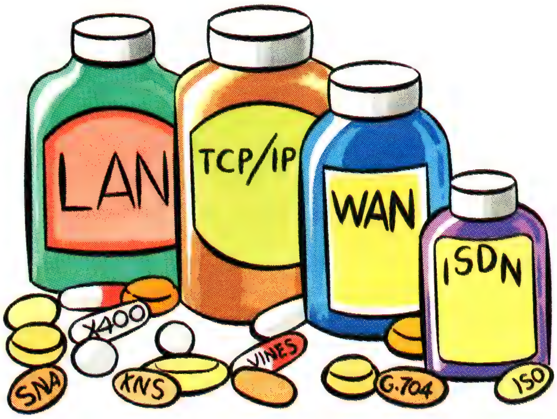

HELP YOURSELF!

To the best range of communications Diagnostic and Management Systems. Information Network Solutions have one of the largest ranges of instruments and network management systems available that will put you back in control. Tools to help even the most complex, multivendor networks. Test and management system for:

• LAN-WAN • ETHERNET • TOKEN RING • X.25 • SNA • ISDN • FRAME RELAY • SMDS • FIBRE OPTIC • MULTIPROTOCOL • MULTIVENDOR

We'll help you find the problem quickly, saving you time, money and your sanity. Don't go looking for network trouble with anything less. Call now on Phone: **(02) 906 6335** Fax: **(02) 906 6327**. Or simply fill in the coupon and return to us today.

INS

Information Network Solutions
Unit 3a, 3 Lanceley Place, Artarmon,
NSW 2064. Fax: (02) 906 6327
Phone: Sydney. (02) 906 6335.
Melbourne: (03) 653 9480

Please send me more information on:

☐ Fibre Optic Test Equipment
☐ Network Management Systems
☐ Coax and Twinax Analyser
☐ Network Aseet Management
☐ PC Based Analyser
☐ Digital Patch System
☐ LAN-WAN Analyser
☐ Telephone Network Simulator
☐ Expert Sniffer Analyser
☐ RMON Based Management System

Name: . . . . Phone: . . . .

Company . . . .

Address: . . . .

. . . . P'code: . . . .

identify themselves before using resources and to authenticate their IDs using passwords. It must also prevent password data from being accessed by an unauthorised user;

- Audit — The operating system must record all access transactions between users and the resources being protected. Information to be recorded includes files and programs that are used, the users that request access, and the times at which access is requested. IDs of network nodes used in any access also must be recorded.

Of the three vendors that have applied for the C2 rating, only Novell has put in for end-to-end approval. That means it's looking for certification that covers the operating system not only on the server side but also on the client side.

"Everyone else has been getting individual boxes approved, but we are trying to get an entire network certified," says Carl Allen, Product Marketing Manager for security at Novell. Novell is hoping to get a C2 rating either in the fourth quarter of this year or the first quarter of 1995.

While Microsoft is looking to have both the server and client components of Windows NT tested, the two software packages will be tested independently rather than in a networked configuration. Microsoft says it hopes to obtain C2 certification as early as September of this year.

Banyan, which is applying to get C2 approval for its server software only, says it is still as much as two years away from being awarded a C2 rating.

Some of Novell's competitors question whether it is even possible to get an end-to-end approval that will cover all parts of an installed LAN. The naysayers point out that the only products granted C2 status so far are standalone computer platforms like the OpenVMS VAX Version 6.0 operating system from Digital Equipment Corporation.

The NSA, however, says that C2 approval for an entire NOS is at least theoretically possible. Possible, maybe — but certainly not easy. Like Novell, Banyan also began by seeking an end-to-end C2 approval for its operating system. It recently abandoned the process because of the difficulties involved in getting the DOS version of its client software certified. "DOS does not have enough security," says Shannon Lewis, Product Manager for Security at Banyan.

Banyan's decision points out a general and chronic weakness in LAN operating system security. The latest operating systems are slanted toward providing impregnable security for the files and programs that are stored on the server. But they provide far less protection for data stored on user workstations. For example, the server part of IBM's LAN Server comes with a feature called the Secure Shell that prevents would-be trespassers from disabling security facilities by tampering with AUTOEXEC.BAT and CONFIG.SYS files. But that feature is not included in LAN Server's client software.

Secure Alliance

Novell has an answer for the sceptics: it says it is planning to boost the security capabilities of its client software to C2 levels via an alliance with security vendor Cordent. Installing Cordent's Assure product, which comprises a board and software and runs on DOS or Windows workstations, will bolster security to C2 levels for NetWare clients, says Novell. "Having C2 certification will cut down on the need for independent testing of NetWare LANs," says Terese Weipert, Vice President of the Information Security Division at Cordent.

Whether corporate network managers are worried enough about potential security breaches to pay up for C2-level security remains to be seen. But the message from some observers is clear: let the non-buyer beware.

Stephen Saunders

DiMONS
Distributed Network Management Platform with Features:

Discovery
With NetLabs/Discovery, you can easily locate devices on a Network and map out Network type, devices and topology. Discovery creates a graphical representation of the Network, using device-specific icons.

AssetManager
First application to automate the time consuming process of finding and tracking your valuable hardware and software resources.

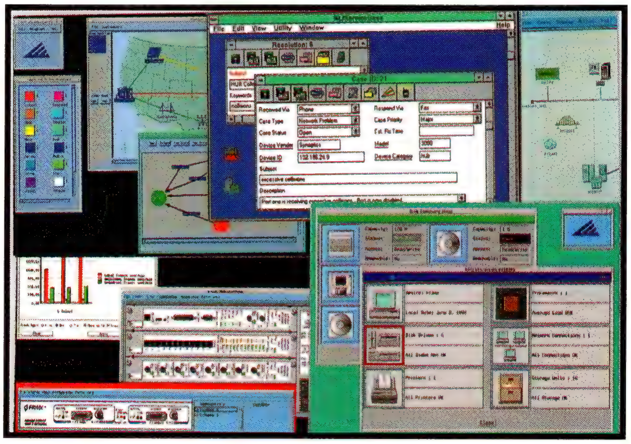

Assist
Monitor Network Performance and manage faults. The MetaViews application provides you with concise summaries and reports of alarm statistics and traffic, so you can quickly identify Network trends and problem areas.

Service Desk
This sophisticated automatic trouble ticketing and help-desk system lets you easily track, route and resolve network, system and user problems through one centralised facility.

The first fully distributed Network Management Platform, NetLabs/DiMONS allows multiple users to manage an entire network from any point on the network. Unlike conventional network Management architectures that depend on one central hardware system, NetLabs/DiMONS frees the users and applications-the programs that solve specific network problems-from system-specific use. DiMONS conform to leading industry standards and integrate applications smoothly with existing and future standards-based solutions including SNMP, OSI, CMIP, and OMNIPoint specification.

For more information, call (Sydney) 906 6335 or (Melbourne) 653 9480.

INS
Information Network Solutions
Unit 3a, 3 Lanceley Place, Artarmon, NSW 2064. Fax: (02) 906 6327

Say hello to your new secretary.

Leave your chair at the office for a meeting or sales call, and you're walking away from potential new business. (Who wants to leave a message when what they want are answers?)

The Electronic Secretary from NetComm changes all that. Answering and transferring your calls. Forwarding messages and automatically faxing out information on request.

It offers a comprehensive, personalised service to keep your office running even when you're not there. It's ideal for the power user who needs to keep in touch and for small offices, home offices or company departments.

The Electronic Secretary connects to your PC (IBM compatible) and includes NetComm's Cooee voice, fax and data software for Windows. It's simple to program via your PC and personalise with your voice. (And it's a fully featured 14,400 bps data and fax modem.) Setting up is straightforward. Electronic Secretaries range from $599 to $1499 RRP. Find out more by giving us a call. We'll be happy to introduce you.

NetComm®

Electronic SECRETARY

KEEPING YOU EVEN FURTHER AHEAD.

Sydney Phone: 888 5533 Fax: 805 0254 Melbourne Phone: 696 2222 Fax: 696 6017 Brisbane Phone: 852 1666 Fax: 252 2255
Perth Phone: 470 5111 Fax: 470 5122 Adelaide Phone: 376 1475 Fax: 376 1475

Words At Work NET 031

LAN VIEWS

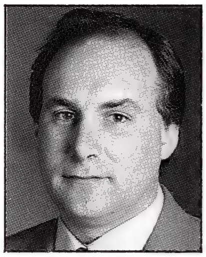

Bottles With Wider Necks

Graeme Le Roux

One of the architectural eccentricities of the PC has been the retention of early bus technology. While the designers of minicomputers or mainframes introduce a new bus technology specifically designed to meet the demands of each new family's processors and peripherals, those who design today's hot i486 based servers and desktop machines are still haunted by the ghost of the original AT.

The reason for this is that the marketing types who wrote the specifications for the hot i486 you are buying today demanded direct bus compatibility with the machine you are replacing. The 8-bit serial parallel card you bought with your AT, which works in the 33 MHz 80386 ISA machine you have now, will work in the 486DX2 66 MHz EISA clone you are putting on your desk today. Unfortunately, if you want really fast graphics of the sort you would find on a workstation you will have to pick a machine with a built-in graphics engine, and if you want a PC database server which can run an extraction over a million record database using six keys in under ten minutes you are probably out of luck — all because your new machine is designed to allow you to plug that 8-bit card directly into its data bus where it will run just as inefficiently as it did in the original AT.

In order to use the processing power of post AT processors with at least some efficiency PC designers had to start putting things like memory and graphics sub-systems on proprietary busses which are 'local' to the CPU — hence the term local bus. Unfortunately, disk and network sub-systems are necessarily still stuck on ISA, EISA or MCA busses because, with the debatable exception of VESA, there has been no local bus standard which has found wide acceptance among peripheral manufacturers. The reason you still have that 8-bit serial/parallel card is that you just couldn't resist having those extra ports — even though you don't use them — right? Fortunately with the advent of today's breed of i486 and the dawn of the 'Pentium Era' (Pentium/Alpha/Power PC Era?) the industry has, apparently, seen the light and it would seem is gradually falling in behind the PCI (Peripheral Component Interconnect) Local Bus (PCILB) standard first proposed by Intel in late 1991.

> *". . . I think PCILB will be a necessary criteria for most larger sites within the next year or so."*

Architecturally and philosophically PCILB is similar approach to that taken by minicomputer and mainframe vendors when they introduce a new family. First, you build a processor and its cache memory any way you like on a CPU bus, then you build a device to bridge between this CPU bus, a standard memory farm and a carefully specified standard peripheral bus (PCILB in this case) to which anyone can build an add-in unit. Most importantly, the peripheral bus is independent of the CPU other than that it runs at the same speed as the CPU and has the same width — i.e. 32, 64, etc. bits — as the CPU. Backward compatibility is provided by device(s) on the peripheral bus (in PCILB's case an Expansion Bus Controller) which provides a perfectly standard interface to older technologies — ISA, EISA or MCA.

The two main advantages to this approach are, firstly, that you can build as many Expansion Bus Controllers (EBCs) as you like into a machine and therefore support ISA, EISA, MCA, (or for that matter NuBuss, TurboChannel etc.) in one machine and, secondly, that this approach is scalable; today PCILB peripherals can run at 33 MHz in 64 bit slots, tomorrow's machines can run at, say, 330 MHz in 256 bit slots with an EBC providing 33 MHz/64 Bit compatibility. Today, PCILB makes 132 MBytes/s available to 32-bit cards and twice that to 64-bit devices. This sort of internal bandwidth is emerging just in time for such applications as SQL servers which require high performance disk sub-systems and future high bandwidth network interfaces such as ATM's 620Mbps variation, not to mention high definition real time — i.e. 'full motion' — video replay.

You can get a feel for the sorts of bandwidth numbers which these applications require with some very simple maths. Consider the video example; given an aspect ration of 4:3 the pixel density on a 640 by 480 resolution 13 inch monitor is 5.87 pixels/square mm. To get the same resolution on a 26 inch monitor (about the size of a large domestic colour TV) requires a display resolution of 1,280 by 1,024 pixels. Assuming 16 bit colour gives a frame size of 2.5 MBytes on this 26 inch monitor, full motion video requires delivery of 30 frames per second — 75 MBytes/s. If the user is running some form of editing application this may mean a data stream from disk to memory (at 75 MBytes/s), one to screen (at 75 MBytes/s) and one from the edit application, in memory, also at 75 MBytes/s. Total bandwidth required by the bus 225 MBytes/s worst case.

The network case is even more straight forward. Consider a multi- homed file server supporting a network of 40 ATM workstations. The lower bandwidths proposed for ATM workstations are in the 20 to 25Mbps range. Forty workstations times 20Mbps is 800Mbps worst case at the server. A careful administrator would almost certainly allow for expansion and redundancy so a server configuration which includes two 620Mbps ATM interfaces seems likely — 1,240Mbps. That's a bandwidth of 147.82 MBytes/s plus sufficient extra capacity to carry on other processing required of the server. (By the way, if you are working through these calculations yourself remember that the 'mega' in megabit, as in storage, is two to the power twenty while the 'mega' in megabits/second as in bandwidth, is ten to the power six.)

To return to our would-be ATM server, handling 147 MBytes worth of data implies getting said data off disk at that rate or better and that is a similar — and possibly consequential — problem to the million record data base extraction where you have to read each record from disk, buffer it and then write it somewhere. If you do the calculations you will find that a RAID array of SCSI drives with an appropriate PCILB controller would fit the bill. By increasing the number of SCSI busses on the controller you can design a disk sub-system which can deliver data at any speed which a memory sub-system can read it.

For the sort of reasons I've discussed here I think PCILB will be a necessary criteria for most larger sites within the next year or so. It will certainly be required by the time ATM (or some other high bandwidth standard) enters general use.

Graeme Le Roux is a Director of Moresdawn Pty Ltd (Bundanoon, NSW) and specialises in local area network consulting services.

OPEN SYSTEMS

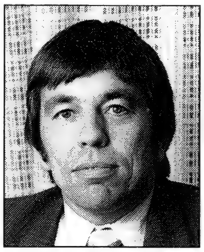

Carrier Technologies to the Desktop

Alan Lloyd

The US FIRP report and recent articles on the TCP/IP and GOSIP debate in this magazine have certainly started people examining the TCP/IP-OSI debate in a closer fashion. In global terms, the situation shows signs of crystallising along the lines of a US vs. The Rest battle with the United States (because it is the controlling agent of IPS technology) believing that it can provide a cost effective IT infrastructure based on a range of multi-protocol environments, while most other countries pursue an OSI strategy. Unfortunately, the cannon fodder may well be users who pick the wrong side.

Those users who sensibly choose to follow the OSI route in designing their networks and plotting their company's future IT directions can be comfortable in the fact that they are choosing the networking technologies as specified by the world standards bodies and telecommunications carriers. WAN and LAN technologies are now converging and the technologies being adopted are derived from the carriers and the ISO/ITU. These days, carriers are not just providing the wide area wires; they are defining large scale distributed systems and how their technologies are migrating toward the desktop. Because of their size and investment in global technologies, these carrier implementations will simply overtake many of the proprietary and de facto standards.

In fact, the old point about TCP/IP replacing OSI is really asking will the Internet group and other proprietary systems replace the carriers, the ISO and ITU? Or will all the RFCs (1,600 or so) replace the thousands and thousands of ITU recommendations and the ISO standards (some of these standards are listed on page 44)? The answer in both cases can only be *no*.

The other point is how many internetworking and routing protocols do we need for true global networking? We certainly need only one form of a global network addressing scheme and that must by definition be internationally defined and administered.

It is widely understood that the evolution of communications and networking technology will take us from leased line and packet mode networks through to ISDN and voice convergence and then on to B-ISDN and ATM switching. This evolution will also embrace interfaces to mobile devices and satellite services.

As we all know, ISDN interfaces and ATM is now appearing in the local environment and at the desktop. ATM into the wide area will take a bit longer to implement because one is dealing with the different parameters such as the cost of national/carrier infrastructure and the need to provide full ATM management and network signalling. The point made here though is that a *carrier* defined technology (I series — ISDN and I.363 — ATM) is now providing the protocols for the LAN environment and direct connections to the WAN. i.e. these technologies are satisfying the requirement for wide and local area bandwidth and cell-based switching.

Internetworking is another contentious area but it's really the TCP/IP and OSI debate in a different guise. CLNP carries OSI NSAPs and/or carrier numbering plans (e.g. E.164, X.121 and Telex codes). CLNP is now available in most routers and is used to support OSI server-based application products such as X.400. However, the move to CLNP in end systems and NOSs is slow simply because one is shaking out the proprietary protocols and NOSs which in certain cases provide their suppliers with a marketing edge. These protocols also lock-in their customers. I often wonder where in fact the NOS vendors are really going. Simply because that once one standardises on the communications, the networking, the file services, the print services, the message and directory services and RPC interfaces, what is really left?

In my opinion, suppliers of proprietary and de facto protocols should now be very wary of using the term 'global networking' with such protocols; simply because these protocols are not truly global. It is amazing how many suppliers say they support OSI and the standards process, yet invent and invest in new protocols. Why? If there are deficiencies in the international standard, all they have to do is just say so and if the case is justified, it will be fixed.

A typical criticism of OSI is that 'the applications are not there.' Let us look at this situation in a different light. I look at OSI and global distributed (application) systems from the X.400, X.500 and X.700 perspective. This is because the carrier-based broadband networks will provide the high performance transmission and networking core of the world's communication system and the messaging and directory technologies will provide the world's message-based application infrastructure — both levels will be managed in the large scale environment by X.700 and TMN-based technologies.

Applications such as Telnet and file transfer (FTP and FTAM) are really point-to-point services and as such do not have the same OSI or global impact as NSAP addressing, X.400, X.500 and X.700. In fact FTP and Telnet over an OSI stack plus CLNP over B-ISDN would be ideal for file and terminal services respectively.

So in this area of distributed OSI applications we now have X.400 and the X.500 products (and X.700 in the future) entering the LAN environment and the desktop.

ISO TP with CCR (X.800) will also emerge later down the track probably in the 3-5 year time frame as the take up for distributed and cascaded transaction services are required. If the MIS manager and the IT planner/network designer saw carrier and OSI-based technologies in this light, hopefully their perceptions on OSI and Open (Distributed) Systems would change. In fact their approach to business/IT system design should also change to accommodate such trends.

Future Networking

It is the conventional LAN/WAN data network that supports the debates on the use of proprietary, de facto or international standards-based networking. How much sense do these debates make in the context of future communications developments; things like multimedia, virtual global networking and mobile communications and the famous information superhighway? (Note to get on the superhighway your 'car' will need a global number plate (e.g. an ISO NSAP). If you have no number, a 'funny number' or a duplicate address, infringement notices will be served!)

From an end system and network interface perspective, some of the H Series of ITU recommendations provide the definitions for audio video (AV), video telephones and AV storage, HDTV data storage, etc. Some of these H Series recommendations have joint text with the MPEG-2 ISO standard. The ISDN (and B-ISDN) recommendations are being directed at carrier and service network integration, while X.300 and other recommendations are directed at data network interworking and the integration of mobile and satellite services. There is a draft set of F.72x series recommenda-

ITU Recommendations

X.1 - X.1xx — Define aspects of digital networks such as classes of service, multi-cast services, DCE and DTE interfaces, packet and circuit mode operation, PADs, multiplexing, transit signalling, network characteristics and QoS, addressing schemes (X.25), etc.

X.200 - X.2xx — Define the OSI reference model, link, network, transport, session, presentation, ROSE, RTS, ACSE protocols and services. The X.200 series covers both the connection-oriented and the connectionless services, network and transport protocol identification and security and the layer management information for the specific layers of OSI. X.200 is the 'OSI' reference model and the protocol stack definitions.

X.300 - X.3xx — Define general principles for interworking and relaying between public networks (of different types) e.g. packet networks and circuit networks or ISDN. Also covers maritime, satellite and mobile services.

X.400 - X.4xx — Cover message handling systems (MHS) and define MTA, UA, MS and access unit protocols and services. X.435 defines Pedi (EDIFACT over X.400) and X.440 defines the message encoding for voice mail. Later additions to the X.400 series cover MHS management in line with X.700 system management and Telecommunications Management Network (TMN, the M series).

X.500 - X.5xx — Cover electronic directory services and define the DSA and DUA protocols and services needed to build distributed directory systems. In addition to the protocols, the X.500 recommendations concentrate on the directory information model and global naming schemes that are necessary to build global directory systems.

X.600 - X.6xx — Cover a range of items such as connection-oriented and connectionless mode services over ISDN, PSTN and PSPDNs, etc., naming and addressing, registration procedures (X.66x), OSI application titles and the OSI abstract syntax notation ASN.1.

X.700 - X.7xx — Cover OSI systems management, protocols and services (CMISE and CMIP), the management functions, management data and the management information model used to define managed objects.

X.800 - X.8xx — Cover OSI security, security frameworks (authentication, access control, non-repudiation, etc.) and the commitment concurrency and recovery (CCR) application protocol used in ISO transaction protocol (ISO - TP X.86x) and remote operations (X.88x).

X.900 - X.9xx — Cover Open Distributed Processing and provide architecture and modelling techniques for distributed system design.

The other major ITU recommendations are the E series, the H series, the I series, the M series and the T series. The E series of recommendations includes E.163/4, the numbering plan for the PSTN, ISDN, etc. It is also used on B-ISDN and Fastpac (802.6). A number of the H series of recommendations cover codecs and audio video (AV) telephone equipment for interworking with narrowband ISDN. The I series of recommendations cover ISDN, frame relay and bearer services, Intelligent Network architecture, Broadband ISDN and ATM. It also covers management requirements for ATM/B-ISDN. The M series of recommendations cover Telecommunications Management Network (TMN) and the maintenance philosophies for managing carrier (public) networks. TMN embraces the use of X.700 protocols (CMIP) and the OSI management functions. The T series of recommendations cover Telematic terminal systems and includes definitions for fax, videotex, OSI session and transport protocols for telematic devices, document transfer coding and the definitions of character sets. Note that many of the above ITU recommendations listed have equivalent ISO documents and share joint text.

---

tions directed at AV and multimedia service specifications over ISDN, PSTN, PSPDN and mobiles. In addition to the above, there is the signalling scheme for ISDN (Q931/B) that supports E.164 and ISO NSAP addresses (as used in CLNP).

In this future scenario the 'end system' is not just focused on being attached or tethered to the LAN. The end system is where a user is and where the services are demanded. These services being administrative, voice, video, image, data or all of the above and in any combination. In this scenario the user and the end system *must* be identified with a global number (an address) and associated with a user identifier. This information can be carried over network dependent or network independent protocols. It follows that debating TCP/IP and OSI is pointless simply because the definitions of network services, networks and end systems are rapidly changing.

From a networking architecture perspective the IN recommendations and TMN recommendations are being used to formulate TINA (Telecommunications Information Network Architecture) which will support a wide range of telecommunications applications such as payphone services, credit card services, 0055 and 13 numbers. These applications are in fact performing the specific user service processing on top of the IN infrastructure which in turn is on top of the carrier network level functions (e.g. PSTN, ISDN, mobile, ATM and X.400).

The carrier networks are evolving in several ways and at different levels. These levels can be referenced at the communications and networking level (ISDN, B-ISDN, radio and ATM); the value added services level (X.400 messaging and X.500 directory services); the architectural level with TINA (TMN and IN) and the (user) service level (mobiles, card systems, call and service management, etc.). The effect of such carrier technologies will cause changes to the customer's IT environment and equipment as compatible protocol and service interfaces will be needed. These interfaces will cover, for example, communications, networking, messaging, directory access, network management and service registration and establishment. The addressing schemes used throughout these interfaces will naturally be the carrier (ISO/ITU) schemes.

The next stage of the process is NOS migration. Some NOS suppliers have already moved to OSI technologies. Database technologies use the industry standard SQL interface (ISO 9075) and user/screen interfaces now adopt Microsoft Windows, MOTIF, etc. At the internetworking level, once the CLNP and ISO/ITU addressing schemes take hold of the end system address points, then the proprietary and de facto components e.g. NOSs, SNMP, FTP and Telnet will have to be transferred to CLNP. (Note it is pointless to have a network based on international addresses and/or global numbering plans that uses a management system that runs on a proprietary protocol.)

Taking this process a step further, once global naming and X.500 directories are used in a broader sense then the network management system will change to using CMIP — because X.500 and CMIP (X.700) use the same name forms.

This migration will take some years, but as I said at the beginning, observe the scope of the carrier technologies and the user demand for coherent global networks, a single point of administration and management, mobility and value added information services. Now take these technologies and user demands and see how they will be implemented across globally distributed information systems. What one does notice is that the carrier, ITU and ISO standards (and their addressing schemes) are slowly but surely falling into place to provide such a system.

In closing I would like to refer to the last few month's articles on the TCP/IP and GOSIP debate. To me there is only one direction the world is going in, and that is toward the carriers and standards-based technologies. The other technologies will survive I suppose but are their days numbered. Where do they fit in the above scenario? Just the user requirements of getting more functions for less cost, having end system mobility, AV teleservices, global addressing and access to intelligent network services, I feel will limit the market for proprietary networking technologies.

I think the 'OSI is dead and slow' bit has now been done to death and readers of these articles should now be questioning those who still say it. OSI is here and more powerful than ever. With OSI you are choosing the networking technologies as specified by the world's standards bodies. You would do this quite naturally with health and safety products so what is wrong with using this approach with distributed IT systems? For example, why would one buy a medicine with 'Request For Comment' stamped on it?

The main issues for the IT strategy planner and network designer are the methodologies used for their IT/system design and how one embraces the emerging carrier services and the evolution to standards-based technologies and distributed information systems.

Alan Lloyd

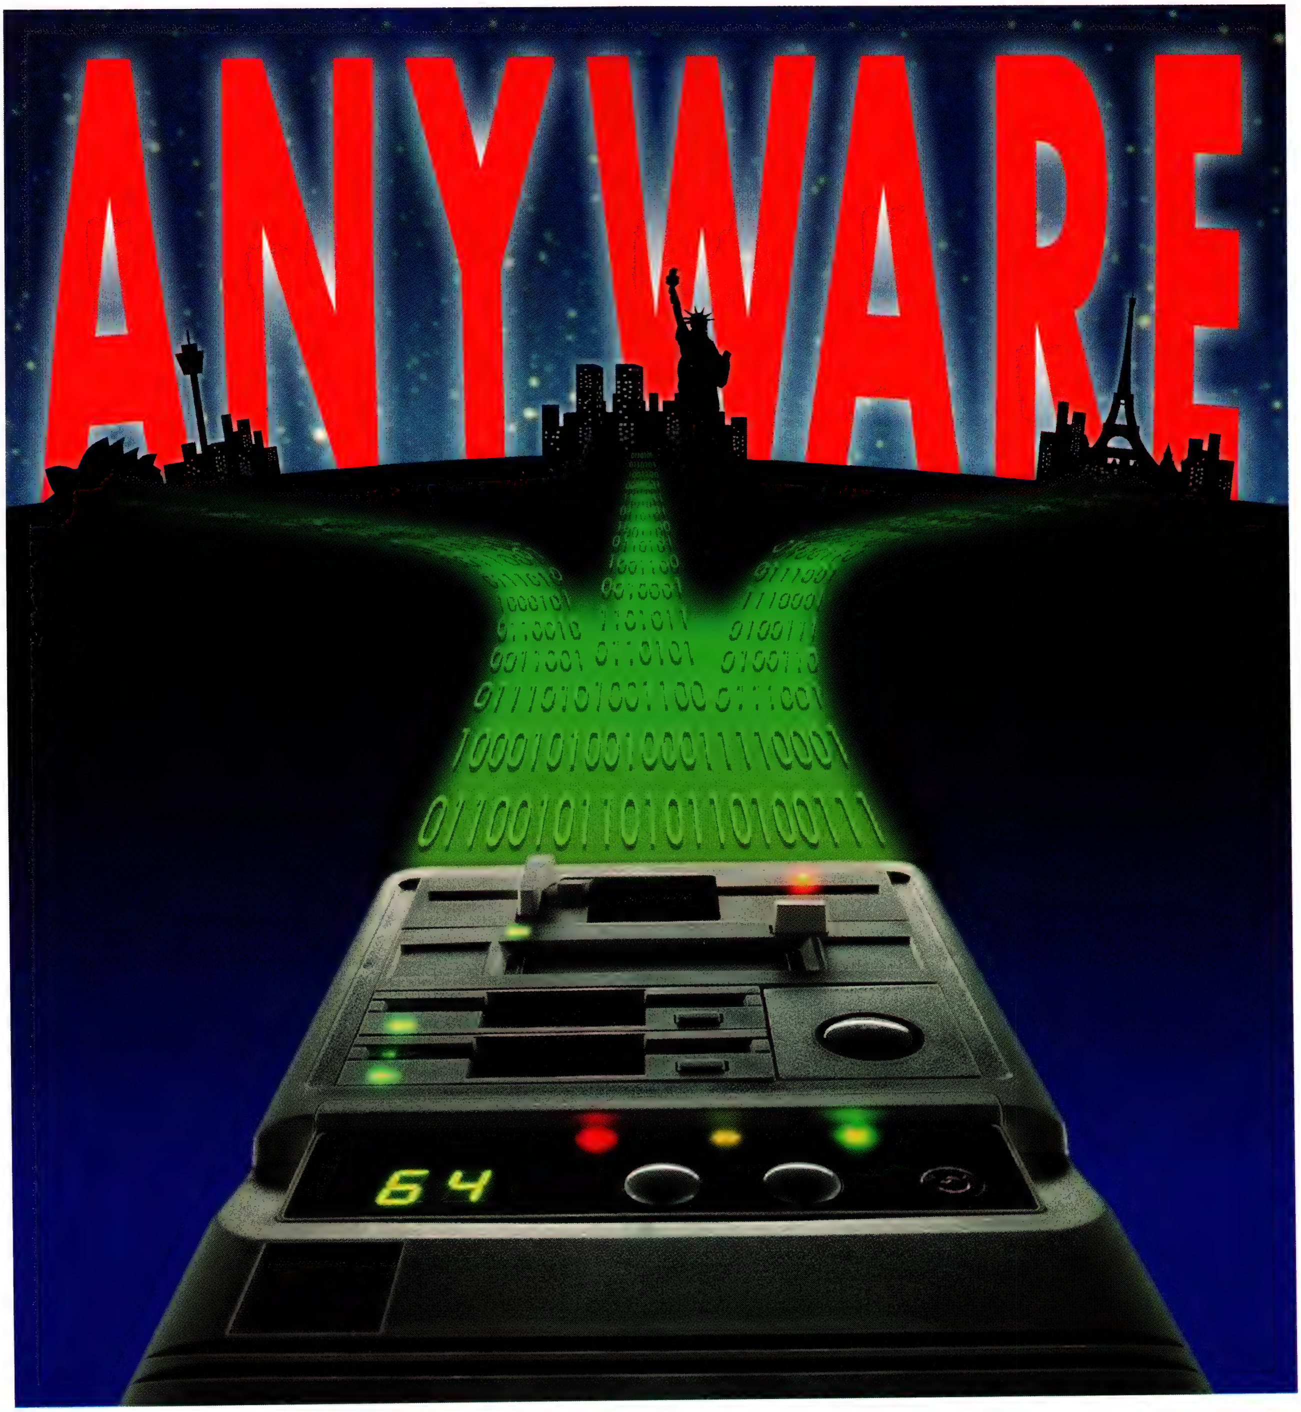

Connect NetWare® with ISIS AnyWare™ and save a fortune on the way!

Connecting your LANs together has never been easier, or so affordable. ISIS AnyWare™ Remote LAN Access products provide on-demand, high performance, transparent and secure connections between NetWare® servers and any other LANs. AnyWare, Anytime, Anyplace!

If you want to know more about our wide range of affordable, reliable, high performance and easy to use Remote LAN Access solutions, contact us for a copy of our *Guide to Remote LAN Access and ISIS*, or talk to one of our product specialists.

Contact our Asia Pacific headquarters in Melbourne on Tel. +61 (03) **629 3800** Fax. (03) 629 5791
or our Sydney sales support organisation, TeleKNX on Tel. +61 (02) **954 5611** Fax. (02) 954 5190

TeleKNX
DISTANCE NETWORKING

ISDN
Powered by ISIS

yes
It runs with NetWare

NETWORK DESIGNERS KNX
LEADERS IN REMOTE LAN ACCESS

Concept, design & imaging: Multicomm Productions +61 (03) 500 9907

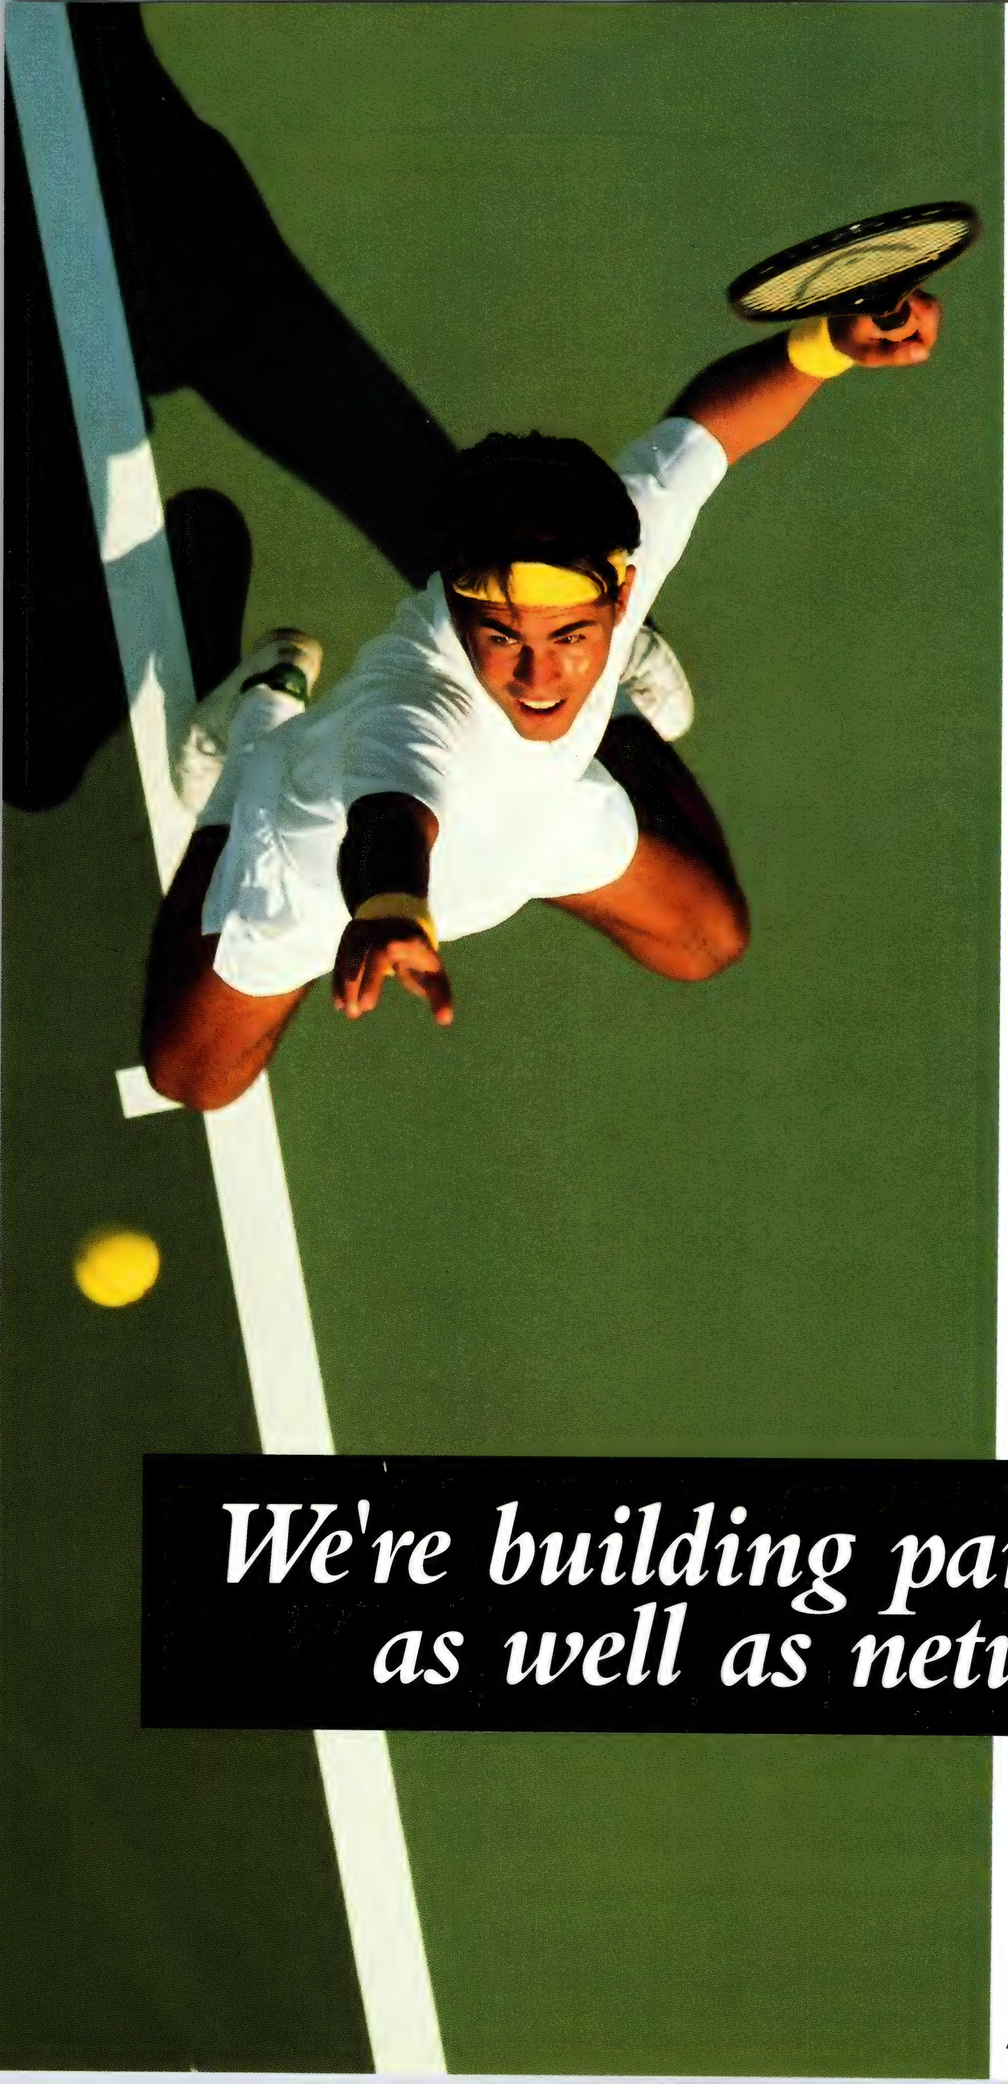

“We can sell the Tennis Centre (16,000 capacity) in about 20 minutes.”

Manager Information Systems
Bass Victoria Ticketing

Making the business of the box office run smoothly across the State of Victoria is what JNA is achieving for Bass Ticketing.

JNA Network Services gets to know your business first, and your business goals, before extending the network. The technical expertise follows. And that is considerable.

For over 30 years we have implemented complex inter-networking solutions for banks, retail chains and government, city-to-city, building-to-building, floor-to-floor. And we get it right.

So see us first. It is quite likely you need more than a network integrator; you need a business partner. Call us now on: Sydney (02) 417 6177; Canberra (06) 273 5044; Melbourne (03)532 8555; Brisbane (07) 870 3434.

We're building partnerships as well as networks.

JNA Network Services

Innovators in Business Connectivity

A division of JNA Telecommunications Limited A.C.N. 005 438 643

J644

Travelling on Magellan's Passport

Nortel's switch combines frame and cell-based trunking to support voice, video and data.

Traditionally there has been a reasonably clear distinction between a company's 'comms,' 'network' and 'MIS' groups. The demarcation line has been usually drawn at the rack housing the modems, multiplexers, PABX interfaces, and so on for which the comms group is responsible and which may be shared by the bridges and routers for which the network group is responsible. This rack(s) is attached to the MIS group's area of responsibility via an HDLC link or the like. Quite often there were separate leased lines carrying voice, MIS data and/or network traffic to remote sites.

This sort of set up is expensive, difficult to manage — in every sense — and above all difficult to troubleshoot. Many IT vendors have seen this, quite correctly, as a market opportunity and as a result this traditional demarcation has blurred; routers have appeared which dispense with the need for separate ISDN terminal adaptors and in some cases modems, MIS traffic has been routed through networks via gateways and various other forms of devices and products like Jtec's J1000 series and Scitec's FastLane have combined many traditionally separate telecommunications functions.

In many cases it is the advent of intelligent telecommunications services such as ISDN and frame relay which have made this process of convergence possible and, of course, ATM is intended to continue the process. Canadian-based Nortel, which is better known for its exchange equipment, has been quietly taking a share of this market with its Magellan range of products, which includes the DPN-100 intelligent Source Routing Bridge, newer releases such as the Magellan Passport 160, and announced products such as the Magellan Access Switch and the Concorde 40Gbps ATM based Multimedia Switch.

An important member of the Magellan family of products, the Passport 160, is a 16-slot, rack mounted, cell-based switch which uses combined frame and cell-based trunking to efficiently support a diverse range of voice, video and data network applications over a variety of local and wide area media standards. Nortel has recently announced the Passport 50, a 5-slot version of the Passport 160 aimed at smaller sites where the larger product's versatility is desirable, but not its physical size, or when its cost cannot be justified.

The Passport 50 may be mounted either vertically in a floor mounting kit or horizontally in a rack. As might be expected from designers working for a company which makes the bulk of its profit from exchange equipment, a great deal of attention has been paid to such things as ventilation and cable management — a detail to which many data networking vendors pay little attention.

As with the Passport 160, the new unit provides frame and cell-based transport with both connectionless and path-oriented routing and is fully interoperable with the other members of the Magellan family. Both Passport models provide frame relay UNI and NNI services at speeds of up to 45Mbps and will support future fibre interfaces at 150Mbps. Trunking from DSO is also provided at up to 45Mbps. The units also provide data transport in the form of cell switching of bit transparent connections for video and other continuous bit-rate traffic and transparent HDLC connections with suppression for HDLC idle flags. Voice support is cell based and includes features such as CAS and CCS signalling, ADPCM Compression at 32, 24 and 16Kbps, speech activity detection, integral dynamic echo cancellation and fax/modem detection.

Multiprotocol Support

At the end of this year Nortel plans to introduce a large range of new data networking features for the Passports including support for routing of TCP/IP, DECnet Phase IV, IPX, XNS and AppleTalk Phase II protocols, along with support for the IEEE 802.1d spanning tree algorithm, transparent bridging, SRB and SRT, remote bridging, and both translation and encapsulation bridging of FDDI. Routing protocols to be supported are OSPF, RIP, Apple's RTMP and EGP, while supported link layer protocols will include PPP, frame relay, X.25 and SMDS.

PRODUCT SUMMARY

Name: Magellan Passport 50

Description: 5-Slot version of the Passport 160 cell-based switch which uses combined frame and cell-based trunking to efficiently support a diverse range of voice, video and data applications

Price: (RRP ex-tax) From $28,000 to $70,000, depending upon configuration

Vendor: Nortel Australia, 495 Victoria Ave, Chatswood NSW 2067 Tel: (02) 325 5200

The Magellan Passport can currently be configured with a wide variety of physical interfaces such as V.35, V.11, DS1, E1, DS3 and HSSI for the wide area, and Nortel plans to make FDDI, 10Base-T, Ethernet AUI and Token Ring interfaces for LAN interconnection available at the end of the year. SONET/SDH (OC-3/STM-3) interfaces will also be supported. Both the Passport 160 and 50 have dual load sharing 800Mbps backplanes and are intended to support future ATM switch modules.

The Passport 160 has N+1 redundant power supplies, optional redundant control processors and optional sparing of DS1 and E1 processor cards, and Nortel plans to release equivalent options for the Passport 50 at the end of the year. The entire Magellan family is manageable via Nortel's NMS and third party SNMP-compatible devices are supported from the workstation-based NMS. The Passport products will be fully SNMP manageable with a future release of software scheduled for the end of the year.

Nortel's Magellan Passports are well thought-out and, in their clean, simple and robust design, are clearly products of a vendor whose bread-and-butter consists of producing highly reliable, easy to maintain devices. The Magellan Passports, with their built-in provision for future upgrades, are designed to have a long service life. As Nortel continues to develop the Magellan range it is possible that the sails of Magellan's ships will darken the horizons of many traditional hub and router vendors.

Graeme Le Roux

A Remote Router With Flexibility

Wellfleet's Access Node provides users at branch offices and other remote locations with a range of routing options.

The days when small branch offices were equipped with a couple of PCs running the same protocol are over. Today's typical remote outpost is considerably more complex, often comprising a mix of Macs and PCs, with a number of users wanting access to the Internet.

Wellfleet, well known for its full-featured multiprotocol enterprise routers, has homed in on the needs of small sites. It's hit the branch office market with an entry-level box that contains the full protocol stack and capabilities of its bigger siblings.

Its Access Node branch office box offers all the functions found in Wellfleet's enterprise routers. It can send traffic to other sites using an array of wide area transmission methods, including frame relay. And the router complies with PPP, so it's not limited to working only with other Wellfleet routers. The branch office box differs from its enterprise counterparts in one significant way. The remote router dials up the central site to obtain its configuration file, eliminating the expense of sending out tech support.

Following Protocol

In its first release, the Access Node (AN) routes TCP/IP, IPX and AppleTalk packets and performs transparent bridging for all other protocols. But the vendor plans to release a free software upgrade that will add routing for OSI, Banyan VINES, DECnet Phase IV and XNS, plus the ability to perform source route and translation bridging.

On the WAN side, the router has a frame relay interface and handles PPP plus its own version of HDLC. The upgrade will add support for DLSw (Data Link Switching) and SMDS, and will deliver ATM capabilities via a DXI (data exchange interface). Looking ahead, X.25 capability is scheduled for September. The AN is priced at $8,216.

For sites firmly set on one LAN protocol, Wellfleet offers two pared-down versions of the AN router with prices to match. The AN Basic Access routes TCP/IP, and on the WAN side it has a frame relay interface and handles PPP plus its version of HDLC. The entry-level router is priced at $5,554.

Wellfleet also sells a scaled-back box with an additional wide area protocol. The AN IP Backbone Software router contains the capabilities of the AN Basic Access and also supports the EGP/BGP router-to-router protocol. It's priced at $6,038. Customers who buy one of Wellfleet's entry-level boxes can upgrade to expand their wide area options, or to include more LAN protocols.

Both of the Access Node's two synchronous WAN ports can accept either dedicated or dial-up links. The LAN port accommodates Ethernet networks and comes equipped with both AUI and 10Base-T interfaces. Wellfleet says a Token Ring version is slated to ship this month.

The synchronous WAN ports operate at speeds to 2.048Mbps and are supplied with EIA-232, RS-422, V.35 and X.21 physical interfaces. Most branch offices will likely use one of the WAN ports for leased-line service and use the other port for less costly dial-up services. The AN's automatic dial-up capabilities give small sites a backup strategy without the need to install a second leased line.

For instance, if the primary line is buckling under the pressure of peak traffic, the router will automatically dial a connection to handle the overflow. When traffic subsides, the AN automatically disconnects the extra link.

PRODUCT SUMMARY

Name: Wellfleet Access Node

Description: A range of full-featured multiprotocol PPP-compliant bridge/routers that support a variety of protocols

Price: Access Node: $8,216; Access Node Basic Access $5,554; Access Node IP Backbone Software $6,038. (All prices RRP including tax).

Vendor: Wellfleet Communications,
Level 4, 60 Miller St, North Sydney, NSW
2060 Tel: (02) 955 2955

The router can also establish a dial-up connection to carry traffic when the primary pipe goes down. When leased-line service is restored, the AN disconnects the secondary link. Offices that need only limited daily connection time with other sites can avoid leased-line charges altogether by relying solely on dial-up circuits.

Administrators can configure the AN to establish dial-up connections at preselected times of day, or whenever data destined for other sites is present. When communications are complete, the router automatically disconnects the link. Users can also set up a dial-up call manually.

Don't Get Choked Up

Users relying on public frame relay services to get high volumes of data to other sites run the risk of having some of that traffic dumped if they exceed their CIRs (committed information rates). The Access Node router has two drawing cards for users concerned about getting frame relay traffic to the other end of the connection intact.

Protocol prioritisation lets users assign one of four priority levels to outgoing data. The highest assures speedy delivery, while the lowest marks frames with a DE (discard eligibility) bit. The DE bit is a signal to frame relay switches that the marked frames are safe to dump in the event that the sender surpasses its CIR.

A congestion-control feature counts the number of FECN (forward explicit congestion notification) bits and BECN (backward explicit congestion notification bits received by the router. Once a preset threshold is exceeded, the router temporarily buffers excess packets until the connection is clear.

Kieran Taylor

Madge-ic Token Ring Adaptor

Madge's new Smart 16/4 AT Ringnode is particularly suited to NetWare Token Ring environments.

The Netherlands may seem a strange place for an alternative to IBM's Token Ring adaptors to be developed, but Madge International, whose parent company is based there, has managed to make a name for itself in just this market.

The firm's biggest customers are those who use Token Ring in conjunction with NetWare and/or TCP/IP, although they also have sales in sites which use IBM's LAN Server. Madge intends to continue its success with the release of its Smart 16/4 AT Plus Ringnode and Release 4.2 of its Smart LAN Support Software. Madge claims that the card 'is the fastest, easiest to install and most versatile Token Ring adaptor for ISA PC applications' on the basis of tests by LANQuest against cards from SMC, Proteon and Olicom, along with IBM's TR II and ISA adaptors under NetWare 3.11.

Madge's new card is a 16-bit ISA adaptor which is EISA-aware. The card uses the TI Token Ring chipset rather than IBM's TROPIC equivalent. The card has not been tested against 3Com's Tokenlink III, which also uses the TROPIC set, or Hypertec's Token Ring adaptors, which have a large share of the Australian market.

When in an ISA bus the card uses programmed I/O, with both 8- and 16-bit modes supported, while in an EISA environment it uses a 16-bit bus master mode.

The card provides both DB-9 and RJ-45 sockets, allowing it to support both UTP and STP cable. Madge has taken a rather unique approach to configuration; all the card's hardware parameters are configurable using switches, software supplied on diskette or software which is resident in the card's SmartROM. This latter option is intended for convenience in the installation of diskless workstations. The Smart 16/4 AT Plus Ring node is backward compatible with Madge's Smart Software drivers and Fastmac Plus drivers from Release 3.3 onwards.

The Smart 16/4 AT Plus Ringnode is shipped with Madge's new Smart LAN Support Software release 4.2, which consists of versions of the company's Fastmac, Fastmac Plus and Smart Software which, in various combinations support NetWare 2.x, 3.x and 4.x Servers, ODI clients (both DOS/Windows and OS/2), NDIS drivers for OS/2 Servers and Clients, DOS/Windows and an NDIS 3 driver for Windows NT and its Advanced Server version. Madge's MADGEODI.LAN driver, which lets NT clients connect to NetWare services, is not shipped with the LSS 4.2 package, but may be freely downloaded from the company's BBS.

Those working with Madge's products in a DOS environment and/or who use Novell's LAN Workplace for DOS will appreciate the fact that LSS 4.2 includes what Madge claims is the first implementation of a downloadable TCP/IP protocol stack. The complete protocol stack is implemented as four Smart Modules; IP, UDP, TCP, RARPD. According to Madge its Smart TCP/IP stack takes up just 6K of DOS memory, as opposed to the 40K required for host-based ODI drivers and Novell's own TCP/IP stack.

Further significant savings are claimed in a multiprotocol environment. Smart modules are processed by the Eagle processor built into TI's chipset, which Madge officials point out allows a workstation to dispense with the protocol mappers normally required when host resident drivers are used and thus provides better performance. Madge's press release states that its Fastmac Plus software 'can process more than 30,000 frames per second, reaching 16Mbps wire speed with frames of 48 bytes or more.' Presumably this is when Fastmac is used in conjunction with the Smart 16/4 AT Plus Ringnode. Unfortunately for the users of other NOS's, Madge's Fastmac Plus software supports only NetWare 3.x, 4.x, ODI for DOS/Windows and OS/2.

In summary, Madge's Smart 16/4 AT Plus Ringnode provides a useful option for those who use NetWare over a Token Ring in an application environment where DOS memory is important — that is, DOS or Windows 3.0 or 3.1 running in standard or real mode — at the cost of the simplicity of the driver software provided by the market leading brands. The level of performance claimed for the card when used in conjunction with Madge's LSS 4.2 might also be of benefit in a NetWare server environment, however it is unlikely that the bulk of workstations — particularly those running DOS or Windows 3.0 or 3.1 running in standard or real modes — could or would generate the level of traffic required to justify such performance.

Madge has also announced an ISA FDDI adaptor called the Smart 100 AT Ringnode which, like the Smart 16/4 AT Plus Ringnode, is an ISA card. The Smart 100 AT Ringnode supports both fibre and copper cabling (SDDI with CDDI support via an optional daughterboard) and conforms to the ANSI FDDI specifications. An AMD FDDI chipset is used.

Graeme Le Roux

PRODUCT SUMMARY

Name: Smart 16/4 AT Plus Ringnode

Description: Intelligent ISA Token Ring adaptor and software

Price: Smart 16/4 AT Plus Ringnode (single card + drivers) $1,371; Smart 16/4 AT Plus Ringnode (10 cards + drivers) $11,449; Smart 16/4 AT Plus Ringnode (50 cards + drivers) $54,711. (All RRP ex-tax)

Vendor: Madge International

Distributor: DDP, 17 Atchison St, St Leonards, NSW 2065 Tel: (02) 906 1200

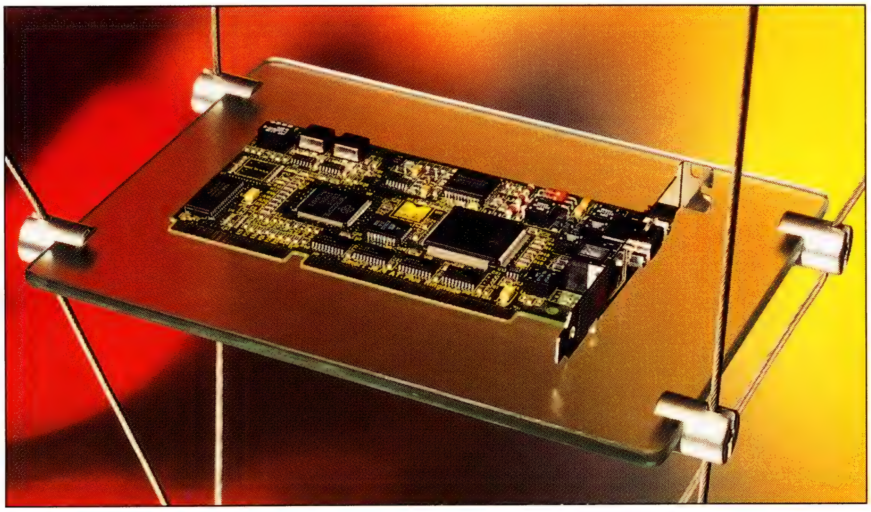

Tony Pool and Malcolm Hamilton
invite you to

MobileCom
EXHIBITION

Australia's 1st Mobile Communications Event

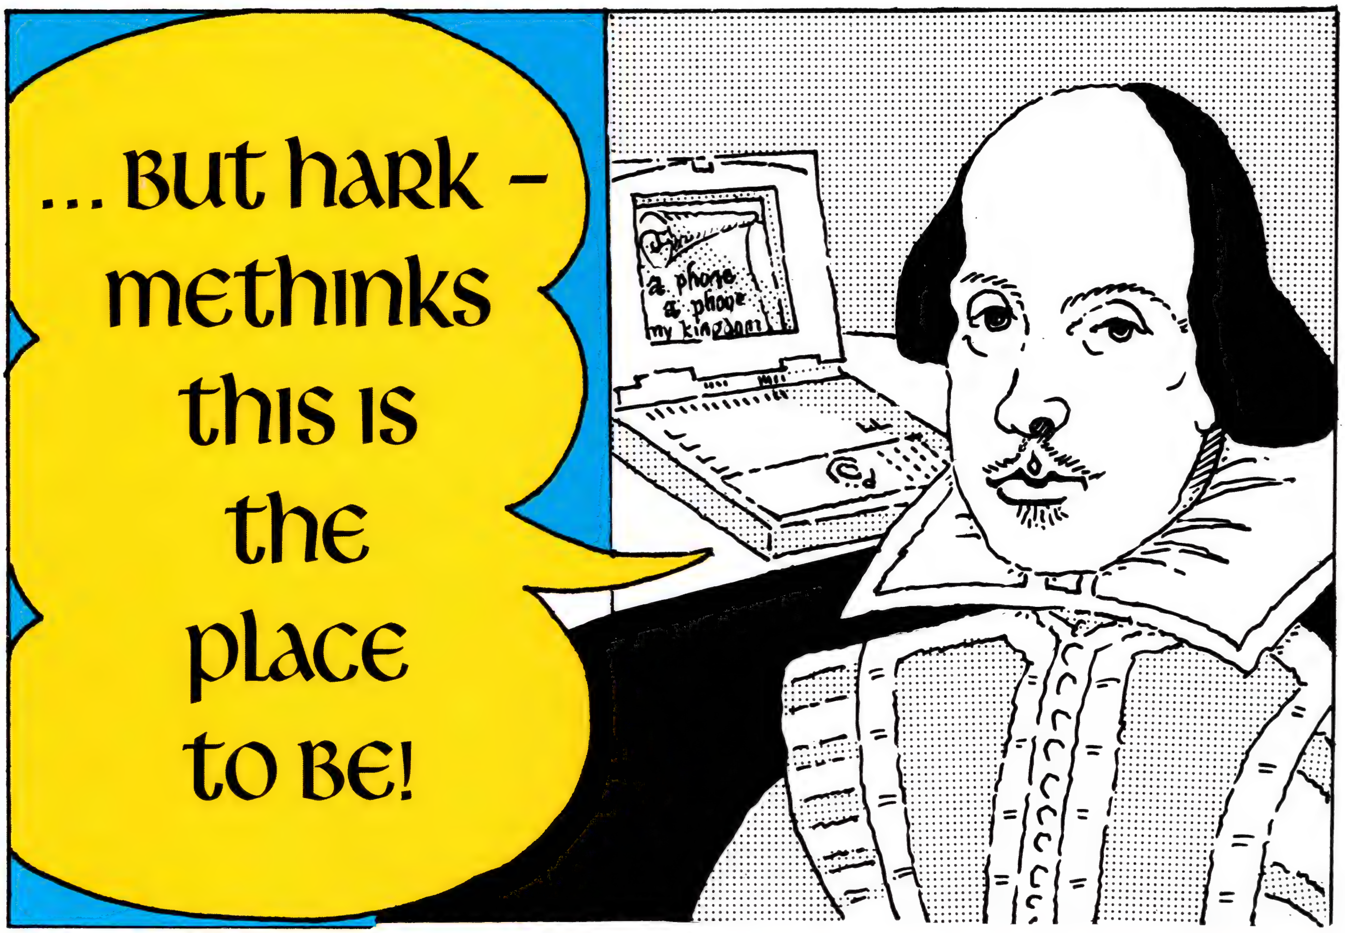
... But hark – methinks this is the place to be!
a phone
a phone
my kingdom

OPEN TO EVERYONE – TRADE AND PUBLIC

SYDNEY	Darling Harbour Friday 25th-Sunday 27th November 1994
MELBOURNE	Royal Exhibition Building Friday 31st March-Sunday 2nd April 1995
BRISBANE	Brisbane Convention and Exhibition Centre Friday 16th-Sunday 18th June 1995

HIGH PROFILE
EXHIBITIONS PTY. LTD.

High Profile Exhibitions Pty. Ltd. A.C.N. 063 398 836 1st Floor, 21 Swan Street, Richmond 3121 Victoria Australia
Phone (03) 428 0415 · (03) 428 0417 · Fax (03) 428 9508 · Mobile (015) 356 698

Fast-Forwarding Telstra's Modernisation

Gerry Moriarty

Managing Director,
Network Products,
Telstra Corporation Ltd

Gerry Moriarty joined Telstra last year as Managing Director, Network Products, with extensive senior management experience in the broadcasting and telecommunications industries. Previously, he held several executive positions in the Television New Zealand Group, including Managing Director of Broadcasting Communications Ltd and Director of Horizon Telecommunications, and from 1987 to 1991 he was Assistant Managing Director of the Australian Broadcasting Corporation. He has a Bachelor of Engineering (Telecommunications) and is a Fellow of the Institution of Engineers Australia and of the Institution of Radio and Electronic Engineers. Liz Fell spoke with him last month at Telstra House in Sydney.

Can we start with the Industry Development Plan which commits Telstra to spending $10 billion on local industry over five years. Do you feel inhibited in any way by lack of capital?
Moriarty: No. I don't believe that is a problem in terms of our current position and current business plans. There are certainly heavy pressures on capital investment to do the major things. The first is to modernise the network, and that means full digitisation, the FMO [Future Mode of Operation] program that was announced last year. Second, we need to invest in systems. There's a substantial amount of work to bring our systems up to an efficient and modern operating standard. Third, we want to increase the range of service, particularly to provide a broadband infrastructure capability.

All of those things, although they are demanding in capital, need to be kept in balance because there is only a certain amount that a company can properly manage at any one time. We are currently able to fund the investment program we have in front of us from our current revenues.

AC: *The 1992-93 results of the Plan show that $3 billion has been spent already. How much do you expect to spend this year?*
Moriarty: It's in the same order.

AC: *That leaves $4 billion over the next three years which would seem to indicate a reduction?*
Moriarty: No. I don't see a falling off. If anything, it may go up, assuming that the profitability of the company is maintained. Five years is a long time in this fast-changing industry and the $10 billion plan probably has a degree of conservatism in the five year projections.

AC: *In relation to the $3.3 billion Future Mode of Operation program, why did you decide to speed up modernisation by about eight years?*
Moriarty: There are significant savings if you're going to address the modernisation of the network very aggressively over a shorter period of time for two reasons.

You can get better value in purchasing much bigger volumes of equipment, and we have achieved that through the contracts that we announced. The second reason we can do it much more cost efficiently is through concentration of construction activity. We can undertake much bigger construction projects and do those more efficiently than a lot of smaller projects spanning a longer period of time.

AC: *Is your area of responsibility limited to the domestic fixed network?*
Moriarty: It's quite complex. My responsibility covers the operation of the domestic fixed network but not the international gateways or MobileNet or the Customer Access Network. I do have responsibility, however, for the planning of the Customer Access Network and therefore for the capital that is invested in it. But I don't 'own' the network. It's not one of my assets and therefore I don't have the operational or the construction staff.

AC: *Surely that means that some synergies and cost savings are lost in this division of responsibilities?*
Moriarty: No. You don't have to own something or have direct line management because of the way in which I operate and the company recognises that it must operate. It's a teaming approach to make sure that we achieve the maximum efficiencies. The FMO in the domestic network, for instance, ran in parallel with a similar FMO in the international network with very close cooperation.

AC: *Was there a separate international contract? I understand there are to be two new international switches.*
Moriarty: Yes. There are two new international switches. But they are being purchased as part of the contract with Alcatel and the SDH [Synchronous Digital Hierarchy] equipment is being bought as part of the Siemens contract. The two planning teams worked together with someone from International sitting on the negotiating team. That's a collegiate way of doing things. I don't have to have responsibility for them. But at the end of the day, in terms of cost, we're talking about something in switching which is less than $100 million in about $2 billion of total contract.

AC: *How do you expect to handle future synergies between the mobile and fixed networks?*
Moriarty: We have a fixed-mobile integration project and have just sent a group, that includes my people and Mobiles, to Sweden to work with a similar group in Telia, the operating company there.

AC: *Telstra's corporate structure was, of course, determined before your arrival. But isn't it more complex to coordinate across divisions. It must multiply the number of meetings?*

Moriarty: No. You'd still have a group of specialist planning people for Mobiles as I have people planning the fixed network.

AC: *Aren't there different budgets for instance?*
Moriarty: At the end of the day, I coordinate the whole of the capital investment budget for all of those areas.

AC: *What will be the shape of the national network when the FMO program is complete? Will there be fewer switching centres, for instance?*
Moriarty: Yes. There'll be some reduction in switching centres and there'll be a reduction in the number of layers of the network. It will be a much simpler network and therefore it will be cheaper to maintain. It will be cheaper to introduce products because it's much cleaner and has less layers. It will be linked by SDH transmission which is an intelligent self-healing technology.

We would expect it to have dynamic traffic routing in the main trunks of the network and so it will use all of the power of intelligence that is available in both transmission and in switching to achieve operational efficiencies and traffic management that is required by a dynamic competitive operating company.

AC: *How many major switching centres will remain?*
Moriarty: The number of big switching centres will reduce to about 200, but they will become bigger so they're handling more lines and therefore we will have more small remote switches attached to them. It's on those bigger nodes in the network where we carry more of the intelligence and therefore we can introduce new products more efficiently. The other thing is that we will reduce the number of transit nodes in the network from about 75 to about 24. Again that's a demonstration of the simplification of the network.

AC: *I assume the recent Network Operations agreement Telstra Corporation has negotiated with the CWU [Communications Workers' Union] means that some staff must go?*
Moriarty: Yes. We have registered an historic agreement in the Commission which is a national milestone in industrial relations and a hugely important achievement for Telstra as a company.

AC: *Why is it a milestone?*
Moriarty: It's historic for a number of reasons. It is essential, for example, to the achievement of our World Best Practice performance objectives. Further, we reached the agreement without industrial stoppage. It's worth noting that it replaces an agreement that resulted from the network outage in the early 1980s.

AC: *But surely the network technology has changed so much since then that the technical operations staff simply can't pull the plug any more?*
Moriarty: The network still requires people to run it.

AC: *How many staff do you expect to go?*
Moriarty: The new agreement provides for the further reduction of the Network Operations staff by up to 1,500 over the next two years.

AC: *That's 1,500 out of how many?*
Moriarty: That's out of about 6,000.

> "The positive thing is that we've achieved an historic [trade union] agreement without industrial stoppage which recognises the massive change that technology is bringing to our network. It recognises that we have to address the implications of that through reduction of staff and it recognises the way we're going to achieve that in conjunction with the union and the staff involved."

AC: *And substantial cost savings are expected to flow from this new agreement with the CWU?*
Moriarty: Yes. There are cost savings agreed with the union as a result of FMO and as a result of a centralisation of the network operations function over the next two years. I must say that I'm delighted we've been able to negotiate that agreement with the principals of the union without industrial stoppage. The agreement itself was put to the full vote of the staff of Network Operations on the 20th of April and was agreed by a majority.

AC: *Did you expect opposition from the CWU?*
Moriarty: Well, we've worked extremely hard on an agreement through working parties over the last 12-15 months.

Since the beginning of this year, it was the subject of extensive industrial negotiation by a team which I led with the principals of the union. The positive thing is that we've achieved an historic agreement without industrial stoppage which recognises the massive change that technology is bringing to our network. It recognises that we have to address the implications of that through reduction of staff and it recognises the way we're going to achieve that in conjunction with the union and the staff involved.

AC: *Will those jobs be replaced in the wider industry? When the FMO contracts were announced, there were claims that new jobs would be created and Siemens referred to hundreds of jobs.*
Moriarty: One never claimed that there would necessarily be a retention of the same number of overall jobs. But there are jobs created in those companies as a result of the manufacturing of the equipment. And in order to implement FMO, there has been a retention of jobs in our Design and Construction Group. We have 7,800 staff and we've retained them.

AC: *But won't some of those jobs go when FMO is implemented?*
Moriarty: The question is what happens in 4-5 years as the FMO program rolls down. We want to look at the other growth initiatives. For example, we're creating 1,000-2,000 new jobs in the construction area using Visionstream to roll out the broadband network.

AC: *That's mainly construction again which is different from technical and design skills.*
Moriarty: That's true. We can't always map these things ideally. But the company has a genuine commitment to finding growth opportunities. I think we've demonstrated that with the broadband network. We've created Visionstream, which is a wholly-owned subsidiary, so that we can do that internally and create jobs. And as part of the Network Operations agreement we've made a commitment to explore with the union all possible growth opportunities and also to handle the reduction of staff by looking within the broader community to identify where there might be employment opportunities. We're very serious about easing that because we recognise that up to 1,500 people in the next two years is a sizeable number.

AC: *Will the cost savings and enhanced functionality of the digital network be passed on to the average user in the form of lower prices?*
Moriarty: The competitive market will drive price reductions. It's demonstrated around the world that where competition is introduced to this and to other industries, then prices to customers generally come down. So we've got to ensure we have a cost structure that will maintain the viability of the business. We've got to address the cost and performance quality of our services and we've got to address the responsiveness to the market so we can keep ahead of the field and continue to introduce differentiated products to our customers.

AC: *When will Ericsson and Alcatel provide full facilities switches in the Intelligent Network (IN) mode so they can offer the same range of services as Nortel?*

Moriarty: We made a decision which recognised that Telecom has developed a very important customer base through Nortel's IN capability, particularly at the upper end of the business market. We have developed that market around that product and capability and we see that as ongoing. But there are other intelligent or enhanced products that sit between that upper end of the business market and the basic telephone service. It's a case of deciding where each of the suppliers can provide the best capability and that's something we will work on with all three of them over time.

AC: *There was some comment in industry circles that Nortel was the main FMO loser. Is that correct?*
Moriarty: We said at the time that the process had been hugely competitive. There was a lot at stake. At the end of the day, Telecom has achieved outstanding value through the mix we announced. That's the important thing.

We will optimise that through the strategic partnerships that we are developing with all three switching suppliers and with Siemens as the transmission supplier.

AC: *Will you be buying more switches from Nortel?*
Moriarty: We have signed a five year agreement and a strategic partnership with Nortel to continue to buy their product.

AC: *What were the major criteria you used to select the FMO switching and transmission partners?*
Moriarty: The tender specification set out extensive requirements including criteria such as the company's track record, the ability to develop product, the time to market, the strength of local support and a whole range of other issues.

AC: *There's a lot of hype about ATM [Asynchronous Transfer Mode]as the key technology for broadband. Was the ability to provide an ATM platform important?*
Moriarty: We looked at each company's strategies for ATM so that we had a good understanding of their technology development. It is true to say that there is no company at present in a network environment providing ATM services. It's an emerging technology and we have a program with each supplier to evaluate their capability as it emerges.

AC: *What's the timeframe for ATM? When do you expect full scale switched broadband capability?*
Moriarty: The environment in which we will be switching broadband within the network is a little way off. We have a group looking at ATM, but I don't think there is a common view around the world on the timing of ATM's widespread introduction to telecom operator networks. We will be implementing several trials in the near future.

AC: *Could you go to other suppliers?*
Moriarty: Depending on how ATM fits into the architecture of the network, that may be a possibility. Each of our three suppliers has interesting approaches to the way in which they can upgrade to ATM. But it's very unclear as to how ATM and broadband switching will emerge in a telecommunications operator's network. It's a technology which lends itself to smaller decentralised switches and therefore we need to think through different architectures.

> *"Each of our three suppliers has interesting approaches to the way in which they can upgrade to ATM. But it's very unclear as to how ATM and broadband switching will emerge in a telecommunications operator's network. It's a technology which lends itself to smaller decentralised switches and therefore we need to think through different architectures."*

AC: *Turning to the cable TV [CATV] network contract awarded to Philips, what is the role of Visionstream of which you're chairman?*
Moriarty: There are two major cost drivers when we look at deploying a broadband infrastructure to the home. One is the equipment, the cable and the systems, and the other is the labour to construct and implement it.

In order to ensure that we got the best value, minimum cost equipment, cable and systems, we used again a very competitive tendering process. That generated a lot of competition because it was a substantial cable TV contract. If you look at cable franchises in the United States, 100,000-200,000 homes is a big franchise. We're going to roll out to 1.1 million homes within three years. On an international scale, therefore, this was a big equipment and network contract with all of the systems, and the bidders competed fiercely. That fierce global competition with a five star field for a sizeable contract resulted in us getting outstanding value.

The second cost driver is construction labour. There were a number of approaches we could have taken. But we also recognised industrial reality and therefore sought the involvement and cooperation of our staff and the CWU who came to the party. So we've negotiated a greenfields industrial agreement to get streamlined work practices.

AC: *It would appear that management has been successful in negotiations with the CWU.*
Moriarty: I believe the union recognises that the company needs to change. They've also recognised that we want to work together to find win-win outcomes for the company and for the union and its staff.

AC: *So how will Visionstream operate?*
Moriarty: Visionstream is very similar to my Design and Construction Group that is responsible for the FMO, the network modernisation construction. In this case, we've got a new entity that will construct the broadband CATV network to the home and a new industrial agreement.

AC: *Can we talk about Telstra's proposed Videostream tariffs which were dropped on my desk . . .*
Moriarty: They're not my territory.

AC: *But the overall business plan must have determined these tariffs in relation to the investment in the CATV network?*
Moriarty: The people that are responsible for the business plan are the CBU [Commercial Business Unit]. They take my costs, just as I provide them with costs to rollout new products on our existing switches or I provide them with costs for new transmission routes. They build the business plans and establish the tariffs.

That's the way in which responsibilities are structured in the company.

AC: *Yet I understand you have talked to cable TV licensees?*
Moriarty: I talk to them to give reassurance about the rollout of the network, the technical and engineering issues. They don't talk primarily with me.

AC: *Do the set-top units fall within your territory? Several company executives have said they are part of the CATV network.*
Moriarty: Correct. We're providing the set-top as part of the CATV rollout.

AC: *So will Videostream staff construct the CATV network and install, maintain and service the set-top units?*
Moriarty: Visionstream is the construction company that will do the lead-ins.

AC: *Sorry. It is fairly confusing with both Videostream and Visionstream. Will Visionstream install the CATV set-top units as well?*
Moriarty: I expect that Visionstream will install these.

AC: *Which company will actually provide the set-top units?*

Moriarty: When we made the CATV announcement, we said that we were going to take a further short period of time to decide on the set-top unit and the associated systems. We haven't made a decision.

AC: *Given the proposed Videostream Premium tariff includes billing subscribers for both the CATV line and set-top unit, will others be able to supply the set-top units?*

Moriarty: It is an integrated system. The set-top box is authorised by the encryption and authorisation system, and the customers are then managed by the subscriber management system.

I expect we will be providing those as part of our carriage service. And we have met the cost of the set-top unit so that it is available for service providers to use which reduces their investment substantially.

AC: *Is Telstra considering financing the service provider, Cable TV Services [CTS], either through direct support or soft entry pricing?*

Moriarty: No. I should also stress that our network will be available to all service providers as a common carrier service and the tariffs will be filed.

AC: *And is Telstra itself planning to deliver a package of basic Videostream services such as the ABC and SBS?*

Moriarty: We haven't made any decision on that as far as I'm aware. That's not to say it won't necessarily occur. I can only repeat that the common carrier infrastructure will be available to all service providers.

AC: *When company executives talk of multimedia opportunities these days, it is difficult to know whether they're talking as a common carrier or as a service provider or both.*

Moriarty: For Telstra, a reason for moving up into those service and application areas is to ensure that we stimulate use of the carriage capability. I think that's a reasonable business direction.

AC: *As opposed to the company itself providing services?*

Moriarty: It depends on how we participate in those areas. We've recognised that we don't have the content skills and so we've sought to work with partners.

AC: *Would you agree with the view that carriers and cable TV operators worldwide are backing away from plans to install fibre-to-the-home?*

Moriarty: I think there has been a lot of hype with people talking loosely about fibre-to-the-home.

Our CATV rollout tender has confirmed very clearly that a hybrid coaxial cable-fibre network is the most cost effective because the cost of the optical to electrical connection is still substantial and we expect it to remain so for some time.

AC: *What is 'some time'?*

Moriarty: There will always be an optical to electrical connection. The question is will it be at the back of the TV set, on the wall of the house, at the kerbside or back at the neighbourhood hub?

For the moment, the architecture that is most cost effective is to have it at the neighbourhood hub and then to reticulate from there in coaxial cable. It's debatable that you will necessarily have to go all the way to the home with fibre to achieve the full interactivity that we see as an ultimate vision. I think it will be a very long time before the capability and capacity of our hybrid network will be overtaken, requiring us to go all the way to the home with optical fibre for services.

> *". . . we have a very short Customer Access Network. Our exchanges are much closer to the customer than in North America, for example. Therefore ADSL technology, which is distance sensitive, has a prospect of working better in our network than in some others. We're going to demonstrate that in the pilot."*

AC: *One US commentator has declared that ADSL [Asymmetrical Digital Subscriber Line] is 'a dying telephone technology.' Has it arrived too late?*

Moriarty: I think it's a brave person who predicts the demise of ADSL until some of us have implemented pilots. Our view is that ADSL has an attractive future, particularly in our network where we have a number of benefits.

By world standards, we generally have pretty good outside plant and equipment, although there are some weaknesses which we are addressing. But most importantly, we have a very short Customer Access Network. Our exchanges are much closer to the customer than in North America, for example. Therefore ADSL technology, which is distance sensitive, has a prospect of working better in our network than in some others. We're going to demonstrate that in the pilot.

AC: *Why has ADSL attracted so much recent criticism?*

Moriarty: I think some of the bad press has come from vested interests including competing network providers who see video distribution on the existing copper wire infrastructure as a threat to their dedicated broadband networks. Our view is the two distribution technologies map very well together. ADSL is a different model. It doesn't bring all the programmes to the gate or to every house. It leaves them back in the exchange and you take one programme at a time to the customer when the customer dials it up. So it's a quite different architecture and it raises quite a number of questions. That's the reason for the pilot.

AC: *And the pilot will require video switching?*

Moriarty: The video switching will take place in the hub which will be in an exchange where we have a video server. If we extend the pilot to a number of coverage areas, then we may put in multiple servers.

AC: *And multiple switches?*

Moriarty: Yes. We haven't decided to go to multiple locations yet. But we may well do that to understand some of the broader architecture issues.

AC: *Your reference to the vested interests of competitors invites a comment on Optus' plans to operate broadband cable to the home.*

Moriarty: Optus has substantial broadband capacity into the home already using its satellites as an alternative distribution technology. But technologies will overlap each other and the cost benefits will change over time. That's going to be the nature of the future. We'll have more complementary and competing technologies as time goes by.

AC: *Do you think residential customers have the dollars to buy all these services?*

Moriarty: There are many companies that have done studies of the pay and subscriber services that will be the services that drive our broadband network initially. I think there's sufficient comfort and evidence in all those studies that there is an established market.

The thing that cable has is that it can bring a huge diversity compared with some other delivery technologies. In the judgement of most people, cable is the long-term technology because it can provide a natural evolution through the early analogue service that we'll provide, followed by digitised services, then full integration between broadband and our existing narrowband services and, ultimately, full interactivity.

AC: *When you were at the ABC did you have that vision?*

Moriarty: I think anyone who has participated in these industries has looked to the long-term and developed this sort of vision. It's exciting and I'm delighted to be in this position at this historic time.

Liz Fell is a freelance journalist based in Paddington (NSW).

Perils of the Dial-Up Generation: Regulating Bulletin Boards

Once the preserve of 'tech-heads,' bulletin boards are now widely accessible to the public and carry everything from the latest freeware to pornography. Melinda Wallman examines some regulation issues.

Australia is the second largest user of bulletin boards after the United States. Figures from 1992 estimate the total number of boards in operation in Australia at between 30,000 and 40,000 — mostly run privately as a hobby. All it takes to access a bulletin board is a modem attached to your computer and the board's telephone number. Upon connection, lists of conference areas and files are displayed on the screen. Conference areas are public discussion areas covering a wide range of topics, and users can move information from their computers to their electronic host and vice versa, to access and contribute to discussions electronically.

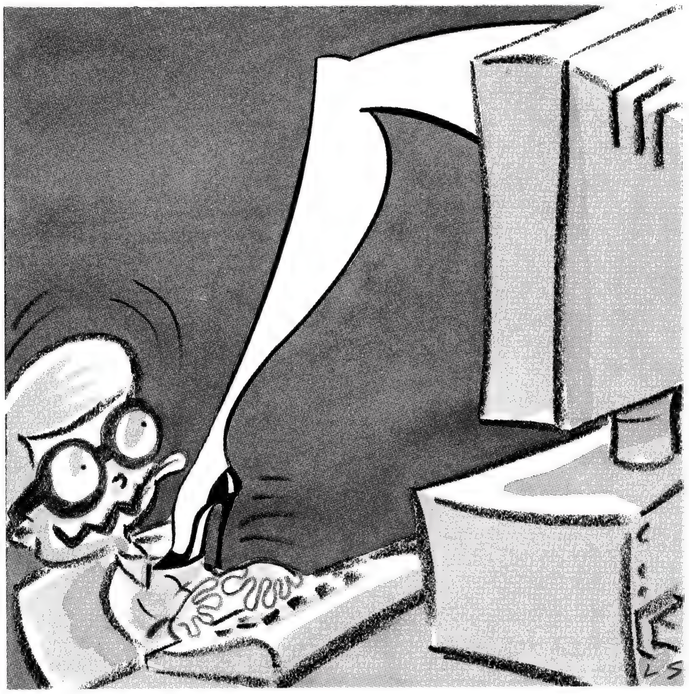

Bulletin boards are administered by a central figure known as the *sysop* (short for Systems Operator). The sysop transfers mail and files between areas, enabling the network to function.

Two modes of communication exist on bulletin boards. Network mail is privately addressed mail communicated via particular bulletin boards. Only the sysop and the addressee have access to this mail. Conference mail comprises all messages uploaded by bulletin board users into a public discussion area. Whether these messages are monitored before becoming publicly accessible depends on the sysop. Some sysops peruse all messages before making them accessible, but most don't have the time to read all messages and only check message areas sporadically.

Board Games

The most common type of bulletin board is the public information bulletin board. These are generally run by dedicated sysops who frequently check conferences to see that messages relate to the relevant topic. If messages are irrelevant or offensive, sysops often delete them. If a bulletin board user continues to upload irrelevant or offensive messages, that user can be cut off or 'toaded' by the sysop.

When accessing public information bulletin boards for the first time, the user must provide the sysop with certain personal information such as, name, address, phone number, age, how the user learnt of the bulletin board and the user's main interest in using the bulletin board. Of course, there are only limited means of verifying the information provided. However, most users of these bulletin boards are *bona fide* communicators whose main aim is to discuss a particular area of interest.

The second type of bulletin board is the commercial bulletin board. Commercial bulletin boards include those run by software companies to assist licensees with application software. Again, to access these bulletin boards users must undergo a registration process. Use of a commercial bulletin board often requires payment of a fee, although some are run as technial support or promotion devices. Bulletin boards can be an effective way to distribute computer shareware or free updates or modifications.

Some commercial bulletin boards contain 'adult' material. When registering to use a commercial bulletin board, the user often has the chance to nominate R- or X-rated areas. Access is usually only provided to an 'adult' bulletin board following verification of the user's age by the sysop, usually by the user sending the sysop a photocopy of his or her driver's licence. In the United States use of sexually oriented bulletin boards is growing rapidly, with the computer generation apparently more comfortable with 'electronic assignations' than oral communications such as telephone sex services.

The third type of bulletin board is the 'pirate' bulletin board, which is traditionally the domain of the computer hacker. As with pirate sound recordings and pirate CD-ROMs, pirate bulletin boards are a real problem area.

State of Play

Most of the current concern with bulletin boards centres on the easy distribution which they provide for pornographic and violent material. There have been recent press reports that a researcher at a UK university had established a bulletin board within the Internet to distribute digitalised pornographic images of children aged as young as two years. The university was notified of the bulletin board by US police who stumbled across it while investigating a paedophile group.

The content of bulletin boards is presently unregulated under Australian law, falling as they do outside the current scope of all censorship legislation. Because bulletin boards are not 'publications' within the meaning of the various State *Obscene and Indecent Publications Acts*, pornographic and violent material can be legally communicated via bulletin boards. Bulletin boards are specifically excluded from the new ACT computer games legislation, the first of its type to be introduced in Australia.

However, when downloaded on to a disk or to print, pornographic material will be caught by obscenity legislation and violent interactive computer games will be caught by the *Classification of Publications (Amendment) Ordinance 1994 (ACT)*.

Even if bulletin boards were deemed to be 'publications' to allow them to come within censorship laws, much like software is considered to be a literary work in order to attract copyright, administering any classification of files would seem a Herculean, if not impossible, task. This is due not only to the fact that there are so many bulletin boards currently operating in Australia, but also to the ephemeral and international nature of bulletin boards.

Because information transmitted electronically into Australia is outside the definition of 'goods' in Section 4 of the *Customs Act (Cth)* the Commonwealth's censorship laws do not apply. If, instead, the material came into the country on a disk through the mail or were carried in a person's luggage the customs legislation would apply. However, amending the definition of goods to include data in a non-material form is unlikely to be of much assistance, as it is unclear how customs officers would be able to detect the electronic exchange of information other than by intercepting telephone lines, which has massive privacy implications.

Federal Task Force

In response to media reports of bulletin boards as breeding grounds for child pornography, drug recipes and bomb-making instructions, a task force was set up in February 1994 under the joint control of Attorney General, Michael Lavarch, and the Minister for Communications and the Arts, Michael Lee, to report by June on how to regulate bulletin boards. The task force is to develop a regulatory approach which would:

- Allow users, parents and guardians to make informed entertainment choices for themselves and those in their care; and
- Provide adequate protection to children from material that might otherwise harm or disturb.

The task force is also to assess whether:

- Existing Commonwealth offences covering the misuse of computers and telecommunications services are adequate to deal with bulletin board abuse;
- Import/export restrictions are capable of controlling the international trafficking and subsequent copying and distribution of otherwise banned material;
- State and Territory laws such as those dealing with the misuse of computers and possession of child pornography are adequate to deal with bulletin boards;
- Current law enforcement tools are adequate to deal with the new technology; and
- Specific offences ought to be enacted, for example, in relation to the use of bulletin boards by paedophiles or for other unlawful purposes.

Regulatory Options

The 'moral panic' generated by the recent media reports has caused concern amongst the many legitimate users of bulletin boards that the Government will adopt a strict licensing system, which would destroy the present ease and fun of bulletin board communication and its many positive functions. While it is important to endeavour to guard against undesirable uses of bulletin boards, the regulatory cure must not be worse than the disease. Bulletin boards are one of the most effective forms of modern communication and are an important commercial avenue for software companies to deliver new product ideas and updates.

Self Control

Given that the number of bulletin boards in Australia runs into the tens of thousands, any licensing system would be very expensive to establish and administer. Also, given that a bulletin board can be readily established by any half competent computer user, policing any licensing system would be a near impossibility. Further, the rigours of a licensing system may encourage more pirate or 'underground' bulletin boards. Either way, a licensing system would soon come to be regarded as 'toothless.'

The only viable option seems to be a system of self-regulation based on a code of practice. The Telephone Information Service Standards Committee (TISSC) structure and Codes of Practice applying to 0055 and 0051 represent an appropriate model for bulletin board regulation. Like the distinction between 0055 and 0051 services by TISSC, a distinction should be made between 'adult' bulletin boards and others. A distinction also needs to be made between network mail and conference mail: network mail, like telephone conversations, should remain free and private.

A formal system of industry-administered registration could be put into place, perhaps using the existing Australian Bulletin Board System Registry. This registry administers the registration of bulletin boards on a voluntary basis and distributes lists of registered systems regularly via both the Internet and FidoNet. Strict registration procedures should be implemented for 'adult' bulletin boards to prevent children from accessing such information. Like the 0051 registration system, this could involve use of a PIN number.

While sysops should have some degree of responsibility for what is communicated via their bulletin boards, if they exercise reasonable care or reasonable precautions in monitoring and censoring of their bulletin boards, then they should not be held liable for illegal acts which take place there.

The code of practice should outline clearly what constitutes 'reasonable care.' The standard should be higher for commercial bulletin boards. 'Reasonable care' must not place so heavy a burden on sysops as to discourage them from operating bulletin boards. In drafting its recommendations, the task force must not lose sight of the fact that bulletin boards are used mainly as a hobby and for academic and research purposes.

Bulletin boards are more 'invisible' than 0055 services, which rely on media advertising. The existence of bulletin boards may only be known by word of mouth or private e-mail. The only effective means of enforcement will be through a complaints-based system where members of the bulletin board community take responsibility for reporting offensive bulletin boards. A bulletin board watchdog (similar to TISSC) comprised of users and Government members could be established to oversee compliance with the code of practice. Complaints of piracy or unlawful distribution of pornography (or other breaches of the code of practice) could be made to that body which could then investigate.

Balanced Outlook

Whatever system is implemented by the Government following the findings of the Bulletin Board Task Force, that system must maintain a balance between freedom of speech, privacy, the right to information, and the efficiency of communication via bulletin boards on the one hand, and the prevention of piracy and the protection of society against the distribution of obscene and indecent material on the other. If Government intervention is to be avoided, the bulletin board community will have to demonstrate a capacity and willingness to pull the plug on aberrant users.

Melinda Wallmam is a solicitor with the technology law firm, Gilbert & Tobin. This column sets out her personal views and not those of the firm or its clients.

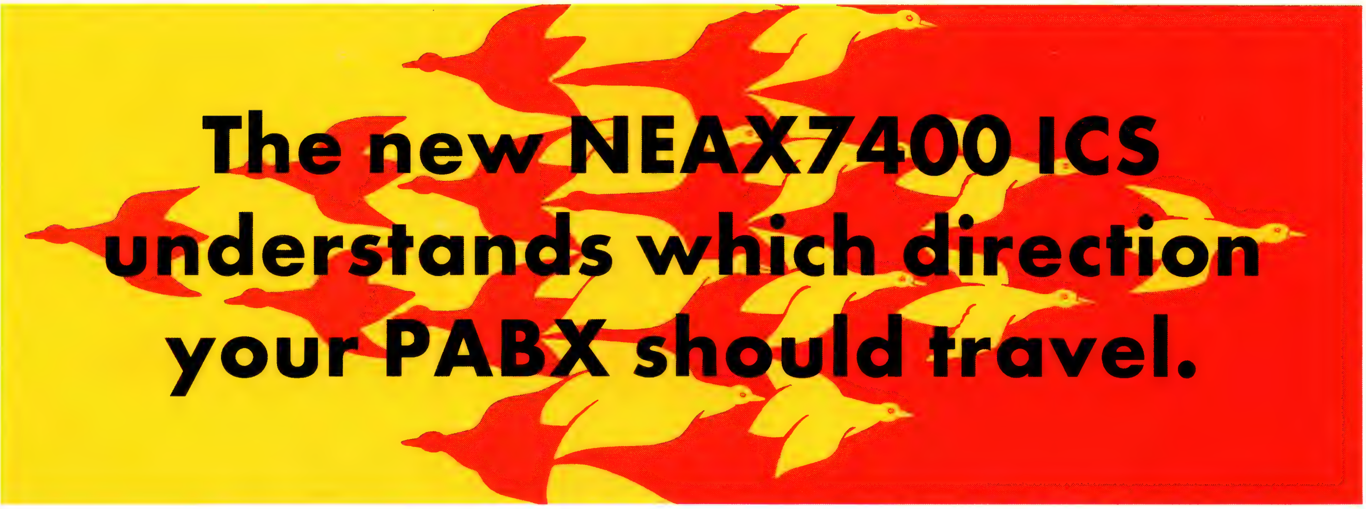

The new NEAX7400 ICS understands which direction your PABX should travel.

The NEAX7400 ICS has been designed as a platform to the future. One which will enable you to integrate your communications with future technologies.

◆

It introduces a number of innovations which enhances NEC's position as a market leader, not only in technology but also in service. There are innovations such as Integrated Voice Compression, the integration of Broadband ISDN and ATM and a sophisticated Network Management System.

◆

The NEAX7400 ICS has the power and the flexibility to meet your current needs. Yet its open architecture environment, ensures the NEAX7400 ICS can be easily updated to integrate new technologies as they become available.

◆

So, to make sure your PABX is able to travel in the right direction and keep up with the future, phone NEC today on Melbourne (03) 262 1111; Sydney (02) 930 2000; Adelaide (08) 301 4900; Perth (09) 220 5555; Brisbane (07) 251 3333; Canberra (06) 250 8777.

C&C for Human Potential ACN 004 803 490 NEC015

NEC

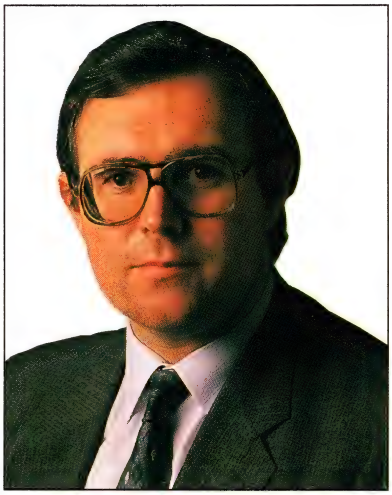

Tom Amos

A Loopy Business

The social aspects of telecommunications sometimes seem lost in the rapid grab for position and services in the lead-up to the open markets of post-1997.

The plight of over one million Australians was brought back into sharp focus by the new Minister for Communications, Michael Lee, in his recent address to the ATUG'94 conference in Melbourne. Lee signalled that the Universal Service Obligations (USOs) had not mysteriously been swept away with the gradual deregulation, but were still real for almost 10% of the community; particularly the unemployed and those on welfare.

It has been traditional wisdom that USOs come into effect in the country areas where the cost of providing a service in remote areas is somewhat higher than in the denser city environment. But the urban areas are where most of the less fortunate users who cannot afford a telephone are concentrated. The question is really whether the price charged for a service is fair and reasonable, given the commercial constraints that GBEs these days face.

One way towards a solution to the social problem would be to consider the basic economics of the local loop, that last bastion of the technology backwater and monopoly. The local loop business is interesting, and the numbers are fascinating. It has been estimated that the average current cost of provision of the next service in urban Australia is near $1,200 per service. This is the cost of the infrastructure to provide a service, all up. There are many parallels overseas to support this basic reticulation provision and, of course, the lifetime of such a service is usually many tens of years.

According to OECD figures, the $1,200 service access cost in Australia has a number of components, including line rental, installation and call costs. For the sake of clarity, add-on services and special features have been omitted from these simple illustrative calculations. It is interesting to note that the line access costs are of course broken into two segments: residential and business. In residential cases the income per access from call revenue is averaged at around $560 per annum, while the call revenue from a business line is near $1,600.

On top of these service provision revenues there is the annual line rental of around $170 and $305 respectively, plus the once-off common installation charge of $120. There are also bond arrangements and so on for new users, on top of these figures.

Now you would expect that the installation charge of $120 would actually cover the average actual cost, and at 10% of the average service provision cost this would seem to be reasonable. Looking at the line rental component in the light of the actual cost of service provision, the residential component has a return of about 14% on the capital invested without depreciation and the cost of money, whilst the business line has a return of over 25% under the same assumption. Both these incomes of course assume no traffic is either generated or terminated in the network.

The ratio of what are long haul to short haul calls of course varies from area to area and service to service, so any allocation for the dispersion of traffic can be only a guide — although the number usually varies between 60% to 70% of the total charge for local traffic. So for residential users the annual bill is about $350 and is due to originating calls, while in the business case almost $1,000 would be a fair average.

So what have we? For the average residential service the simple gross return for the initial investment is over 37% per annum less the cost of the money and maintaining the investment. This would be considered by most standards a good return across all services. For the business services which are identical in technical terms the gross return per annum is closer to three times the residential return under the same arrangements. Who said the local loop business was a loss-making arrangement? Gross returns allow almost the recovery of the original investment within the first year of operation — cellular-like performance!

Even these simple models underline two stark problems which the Minister may have identified as social issues. There seems to be little relationship between the cost of supply and the prices charged for service, which is reflected in the standing charges and the recurring revenue mix. The difference between residential and business is just bewildering, and obviously it is not on any business basis that is commonly understood or publicly supportable. Telephones are obviously a very special product that attract a penalty for bulk use and a disincentive for actually connecting and using them.

Is there scope for retariffing the local loop and, in doing so, to address the problems identified by the Minister at their source, rather than by external Band-Aids that will be open to future iteration and subsequently lost in the jumble?

Many years ago there were claims that there needed to be a rebalance between long haul and short haul to move towards cost-related pricing. It appears that we have well and truly corrected this imbalance, and when you consider that there is a basket-based USO scheme to pick up the country areas that have contributions based upon carrier-timed minutes, then the local loop is really a closed business. Any business with a payback on almost half of its infrastructure in less than a year should be in a position to easily accommodate easy-entry social tariffs rather than punitive arrangements, particularly when over 80% of the infrastructure has already been sunk. It is here that the Minister may find the solution to the social issues in exchange for the status quo. After all, if the line rental does pay for the infrastructure then the calls could be free or vice versa!

The Minister and the numbers bring into clear focus who is getting the real benefit of the massive modernisation of the Telecom plant. It certainly is not the business user, nor the socially disadvantaged end of the equation, so it must be staying within Telecom. The Minister was right, the disadvantaged must be feeling ripped off.

Tom Amos is a partner with telecommunications consulting engineers Amos Aked Swift.

Trish Benson

The Importance of Universal Service

The future shape of the telecommunications industry is a major obsession at the moment. Minister Lee recently stressed "that the regime which is in place after 1997 will be more competitive and liberal than is currently the case . . . underpinning this approach is the strong preference that the Government's role should be minimal to protect consumers and facilitate the development of sustainable competition."

The problem for residential consumers is that there is a strong suspicion that the Government is leaving the market to wave a magic wand and meet all consumers' needs. Major groups of residential consumers have not received the promised benefits of competition in telecommunications. Strong pro-consumer policy and regulation is needed before and after 1997. Universal service — the concept of making telecommunications universally accessible to all — is a case in point.

In its recent report *Reforming Universal Service: The Future of Consumer Access and Equity in Australian Telecommunications*, the Consumers' Telecommunications Network (CTN) argues that universal service comprises five elements:

- Universal geographical availability;
- Universal accessibility;
- Universal affordability;
- Universal technological standard; and
- Universal telecommunications and participation in society.

Universal geographical availability refers to equal access to telecommunications services regardless of one's location. The current Universal Service Obligation in the *Telecommunications Act 1991*, while not an ideal mechanism containing a limited notion of the cost of providing such service, does at least provide reasonable protection for this aspect of universal service.

Universal accessibility relates to the need for universal service to include an access principle. From this perspective, we argue that the term 'standard telephone service' to which the universal service obligation applies in the 1991 Act needs to take explicitly into account a person's actual ability to use the equipment provided. The fact that the current standard telephone service is too narrowly defined has led to inequalities such as the lack of access for people with hearing and speech disabilities who use text communication such as teletypewriters (TTYs) on the telecommunications network.

Universal affordability of telecommunications includes access to and ownership of telecommunications, as well as issues of usage. Changes to the pricing structure of telecommunications ushered in by the duopoly lends urgency to the task of ensuring the universal service obligation includes affordability as a key aspect, otherwise many residential consumers will be disadvantaged at best, or off the telecommunications network at worst.

Incorporating a universal technological standard into the definition of universal service is crucial to allow residential consumers equitable and affordable access to new technologies. There is a need to adopt a dynamic definition of what constitutes a standard telecommunications, or perhaps communications, service, particularly in the face of developments in the intelligent telecommunications network, enhanced network services and broadband technologies.

The current definition of the standard telephone service does not take into account the needs of many consumers and sets universal service at an extremely basic level of technological service. The lack of a clear mechanism for review of the universal service obligation leaves the way open for large business interests to set the level of the standard service without adequate reflection of changing technologies or the needs of residential consumers.

The fifth element of universal service defined by CTN is the concept of universal telecommunications and participation in society. Participation in society increasingly depends on access to telecommunications. A stronger argument can be put that while telecommunications has a dual historical role as both business and social instrument, business use of telecommunications actually rests upon the 'infrastructure,' as it were, of residential access and use.

Measured against this model, the Government's policy on universal service only safeguards one of the five elements of universal service — universal geographical availability of telecommunications. The other four aspects are not effectively addressed. The Government decided in 1991 that these other four aspects would be provided on an ad hoc basis from consolidated revenue or by telecommunications carriers out of corporate largesse. However, little or no government funding has eventuated.

In order to provide an adequate basis for Australian telecommunications into the Year 2000 and beyond, a more comprehensive definition of universal service is needed. This has recently been recognised in the US with the Clinton Administration's National Information Infrastructure initiative.

A comprehensive definition of universal service — including the five elements outlined — needs to be funded and delivered by the telecommunications industry, to avoid the buck-passing of the last four years. In addition, a watchdog and champion for universal service must be established with a clear charter to safeguard all five aspects of universal service — perhaps Austel, though its future role is unclear. Achieving universal service is an urgent task. If the current definition of the Universal Service Obligation is not expanded to meet current needs, future access to the telecommunications network and the new technologies and benefits of competition will not flow to residential consumers.

CTN is the national coalition of consumer and community organisations representing residential telecommunications consumers. Copies of the CTN Report can be obtained by contacting CTN at 8c/245 Chalmers St, Redfern NSW, 2016, Tel: (02) 318 2026, Fax: (02) 318 2031.

Trish Benson is Coordinator of the Consumers' Telecommunications Network, and is based in Sydney, NSW.

EVERYBODY DOES VOICE COMPRESSION, BUT OURS IS WORTH TALKING ABOUT.

FastLane by SCITEC is a new low-cost, multi-function packet bandwidth manager that integrates data, LAN, Group 3 fax and voice on the one leased line, giving it the potential to completely eliminate STD charges and in most applications results in drastic savings that may well pay for the whole system within just one year.

FASTLANE

Designed and produced in Australia, FastLane is the culmination of 150 man years of research and development resulting in breakthrough technology that's left our competition virtually speechless.

When combined with quality manufacturing to ISO9001/AS3901 standards and a design philosophy based on providing true value for money it's little wonder we've come up with the most cost-effective network solution available today. Here is the good news: Frame Relay, Fast Packet Switching, 9.6 kbit/s to 2 Mbit/s, 8:1 voice compression, voice switching and silence suppression, SNMP Management and any '7E' protocol like X.25 and SNA, automatic call re-routing, Ethernet bridging and IP routing plus a migration path to ATM.

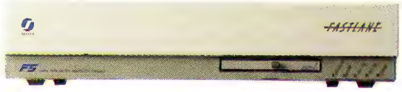

Talk to your local SCITEC office today to hear all the good news for yourself.

SCITEC

NETWORKS THAT PAY FOR THEMSELVES

SYDNEY	MELBOURNE	CANBERRA	ADELAIDE	BRISBANE	PERTH	UNITED KINGDOM
+61 2 428 9555	+61 3 690 8622	+61 6 285 1844	+61 8 362 4000	+61 7 371 9122	+61 9 324 1599	+44 753 833 222

SCITEC PRODUCTS ARE SOLD AND SUPPORTED IN 32 COUNTRIES AROUND THE WORLD

Making the Switch to Ethernet

Switching schemes are breathing new life into Ethernet, the oldest and still the most successful local area network technology. Are they enough for network managers to make do with until ATM arrives?

When Ethernet was first developed it was intended primarily as a way (with the aid of terminal servers) to let a single terminal connect to a number of hosts and as a method of providing a cost effective, high bandwidth connection between hosts at a single site without a rat's nest of serial cables and patch panels all over the place. In this sort of environment it is very hard to overload a simple Ethernet, composed as it is of a number of 10Base-5 segments interconnected by repeaters to which hosts and terminal servers are connected via transceivers, drop cables and AUI ports. As these networks grew, bridges had to be developed to replace the repeaters and allow administrators to contain and control traffic and then provide interconnection with remote sites. Later, as it become desirable to interconnect sites with self-contained administration, to construct firewalls around specific network segments and hosts, and to communicate with other organisations, routers were developed and deployed.

One of the things which made traffic control, and therefore bridges, necessary was the advent of the PC. Users could, and did, start throwing files around the network, rather than just keystrokes — thousands rather than tens of bytes per second. These 'huge' amounts of data were being moved at terminal interface speeds — 9,600bps was high-speed — until Dr Bob Metcalfe and his team at 3Com invented the original EtherLink adaptor and EtherSeries software (Metcalfe, the 'Father of Ethernet' originally developed the technology while he was with Xerox). Users could then start acting as hosts and move hundreds of kilobytes of data around at 10Mbps. Simple Ethernet topographies were quickly overloaded; a fact which is often painfully apparent to the administrators of even relatively small PC-based networks today.

The first and perhaps most obvious solution to the problem was to deploy multiple backbones. As shown in Figure 1 on page 62 you simply run another Ethernet cable, install dual Ethernet interfaces in your hosts and servers (making them dual homed), and dual bridges on each floor of your building. Effectively the network is cut in half and if you install the second backbone in a physically separate location —say a riser on the other side of the building — then you also get some degree of fault tolerance. If one backbone breaks you still have half your network operational. Notice that, while this system has some fault tolerance, it does *not* have any degree of redundancy. Of course as the network grows the original problem of insufficient bandwidth occurs on each backbone, and so the network has to be segmented again. Eventually this course of action leads to a collapsed backbone configuration, as shown in Figure 2 on page 74.

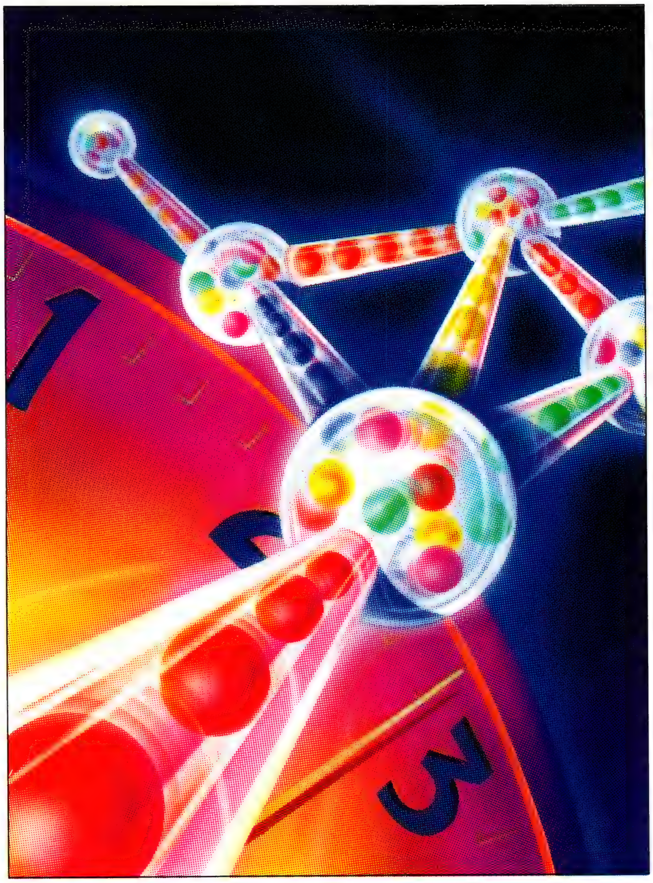

The hubs used in such a collapsed backbone topography typically have a backplane which has a bandwidth of several hundred megabits per second, may be configured with bridges and routers at a lower cost than standalone units, and may permit hosts and servers to be connected via high-speed LAN interfaces such as FDDI. But they cannot remove the basic bottleneck which all Ethernet standards suffer: when one station on a given network segment is transmitting nobody else can. Any attempt to do so results in a

Figure 1: Bridged Ethernet With Dual Backbone

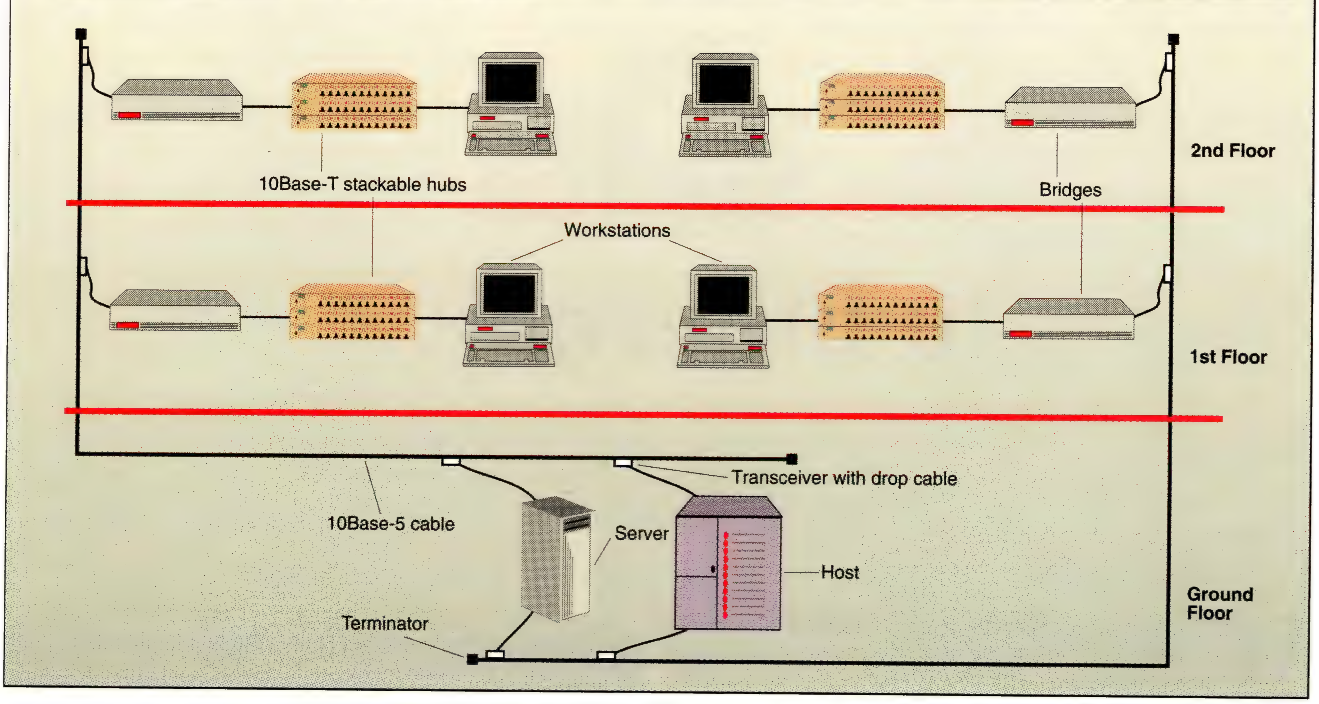

collision and re-transmission. Nor do routers or bridges resolve this problem because, in spite of their speed, they are still logically serial devices. If two stations on one side of a bridge attempt to transmit to separate destinations on the other side of the bridge the bridge will buffer the second packet it receives until it has processed the first. It has to, since Ethernet, like all LAN standards, is a serial medium.

The method by which stations on an Ethernet share this single channel (Ethernet's access scheme) is Carrier Sense Multiple Access with Collision Detection (see 'CSMA/CD and All That' on page 70). CSMA/CD is quick and effective; unfortunately it is also asynchronous and bogs down under a continuously high packet load. It is a high number of stations attempting to transmit frames and not, as is commonly passed on as net-lore, the size of the frames, which reduces the throughput of an Ethernet. Frame size is important to devices, such as bridges and routers, which have to buffer packets, low speed WAN links, and processing devices such as servers and hosts.

The inescapable basic properties of Ethernet dictate that, no matter how many backbones you deploy, or in what form you deploy them, once one station transmits to a given destination over any number of segments — be they bridged, routed or just hanging off repeaters — no other station on any of those segments can transmit. Even 10Base-T LANs in which every station, server and host has its own 10Mbps 'segment' connecting it to a hub (which might have a gigabits-per-second backplane), must obey this rule.

As the number of stations in a network rises, more stations will be forced to wait and the network's response time rises exponentially — and as far as the users are concerned the Ethernet dies.

But what if you could dynamically split your network into small LANs composed of a server, or host, and all the stations which needed to access them at a given moment? What if you could logically reduce the number of stations in each of these LANs to just two, the source and destination of a single frame? This is the goal of all the Ethernet switching schemes which are available; however while the goals of these schemes are equivalent their mechanical details are not.

A Matrix of Bridges

The most straightforward type of Ethernet switching is *packet switching*. Examples of this type of device are the Kalpana EtherSwitch at the low end of the market and the Chipcom (ex-Artel) Galactica at the top end. The simplest way to describe these devices is as a matrix of bridges with a shared address table. This matrix permits any number of stations to transmit at once, provided each station transmits to a separate destination and that the destination is not busy. Just like a telephone exchange using packet, rather than circuit, switch technology.

Consider a small network composed of a Chipcom Galactica packet switching hub with 32 10Base-T ports to which are connected 32 workstations in a simple star topology which are capable of operating as both clients and servers — e.g. Windows for Workgroups, Windows NT, SparcStations running Solaris, or VAXStations running OpenVMS. In this configuration the Galactica will contain four, eight port 10Base-T interface modules, each with its own i960 CPU and a single system management module all of which are interconnected via a 1Gbps capable backplane. Only traffic between ports on separate interface modules is transmitted over the unit's backplane and, since a minimum of two workstations are required to make an Ethernet network, the unit as a whole can physically support a maximum of 16 concurrent 'networks,' so at worst the Galactica, apparently, needs to pass just 160Mbps. So why the 1Gbps back plane? Because, unfortunately, things are not that simple.

Aside from a Cylon attack, the worst that can happen to a Galactica is to have a full size Ethernet frame, followed immediately (i.e. 9.6 microseconds — the minimum Ethernet interpacket delay) by a minimum size frame, arrive simultaneously at each of its 32 ports, with each of those frames having a destination which requires that it be passed across the unit's backplane.

Since the Galactica's interface modules have the same bandwidth as the unit's backplane, the Galactica's internal logic has a consistent 1Gbps bandwidth. Ignoring preamble (7 octets) and the start of frame delimiter (SFD: 1 octet which is always 10101011) the maximum frame size for any packet will be 1,518 Octets (1 Octet = 8 Bits, hereafter Octet and Byte will be used interchangeably unless otherwise stated), thus the Galactica's internal bandwidth is sufficient to handle just over 82,345 such

Continued on page 70

Ethernet Switches Evaluated

Round up the usual suspects. More powerful LAN servers, bandwidth-hungry imaging and graphics, the relentless rise in the number of users on each segment: network managers know all too well why — after more than a decade of reliable service — their 10Mbps Ethernets are sagging.

And they're equally familiar with the technologies typically suggested as solutions: FDDI, fast Ethernet and ATM. But over the past two years a new class of products has emerged that can boost bandwidth on overburdened Ethernets while working with conventional cabling and adaptors. Known as Ethernet switches, these newcomers offer 10Mbps per port and have the potential to multiply Ethernet bandwidth almost indefinitely. For example, a switch with more than 10 ports can furnish greater aggregate bandwidth to users than a 100Mbps FDDI ring. A chassis equipped with a high-speed backplane and a full complement of switching modules can deliver an aggregate bandwidth of more than 1Gbps. And bear in mind that this heretofore unheard of Ethernet throughput comes without having to toss existing equipment on the scrap heap.

But which Ethernet switch is the one to buy? Vendors promise that their products can hit wire speed on all ports — no matter what. To put those claims to the test, the Paris-based European Network Laboratories (ENL) evaluated seven of the first Ethernet switches to make it to market.

ENL, working in conjunction with test equipment vendor Wandel & Goltermann, created the Switch Test Bench (STbench) test suite just for these evaluations (see 'Test Methodology' on page 66 for a complete description of the test suite). Simply put, STbench generates random patterns of multidirectional traffic, creating 1,920 virtual circuits on the ports under test. This level of activity makes it the only test suite capable of stressing Ethernet switches to see how they really stand up under heavy workloads and real-world conditions.

This round of evaluations had several goals. ENL wanted to determine how Ethernet switches can handle heavy loads without dropping packets. This is critical, since these switches are headed for Ethernets where saturation is the No. 1 problem. The test suite also measured latency under real-world conditions to help establish how efficiently switches work under heavy loads. In addition, the tests assessed each switch's ability to deal with errored packets and overloads.

How did the products stack up? Vendor claims to the contrary, only two of the switches could deliver 10Mbps per port under full loads without dropping any frames. The LANswitch from Lannet Data Communications, which comes as a card in the vendor's Multinet hub, was the top performer. Its 1.28Gbps backplane allows it to support 128 dedicated Ethernet ports — the highest number of any chassis tested.

It also was the only switch to guard against frame loss caused by overloads. And it discards some of the most common bad packets found on an Ethernet.

The Etherswitch EPS-1500 from Kalpana achieved the same level of throughput as the Lannet product and exhibited the lowest latency of any device. Its chassis is limited to 15 ports, though, which means that units must be cascaded if more than this number is needed. This can be done via a 10Mbps half-duplex link or with the vendor's proprietary 20Mbps full-duplex connection.

The Powerhub from Alantec also performed impressively — hitting full wire speed on every port when tested with all frame sizes except 64 bytes. In the last case, it dropped only 1% of the traffic. Powerhub also deserves notice because the same chassis accommodates enough FDDI modules to build two rings. And Powerhub not only switches frames but also routes IP, IPX, and DECnet between ports. It screens bad packets as well.

Performance for the other switches fell off — sometimes markedly. The Galactica from Chipcom, ESX-820 from Cabletron, Elite ES/1 from Standard Microsystems Corporation (SMC), and the Linkbuilder 3GH from 3Com all lost more than 10% of small frames when fully loaded.

These results are telling because the majority of packets on the average Ethernet are 64 bytes long, the minimum Ethernet frame size. And when NetWare servers running standard business applications are considered, as much as 75% of Ethernet traffic can consist of 64 byte frames. Bear in mind that retransmitting dropped packets is not only inefficient (because it involves cumbersome higher-level protocols) but also adds to the delays experienced by end-users on an Ethernet.

A Switch in Time

Ethernet switches could well be lifesavers for power users of Unix workstations or scientific and engineering PCs, as well as business users with imaging applications and multimedia aspirations. They also stand to benefit companies where the number of PCs or workstations on a segment is hitting 50 or more — especially if users have made graphics a regular part of network traffic. And any Ethernet site slowed by sickly server performance could get a big boost from these products. What's more, installation is a breeze: all network managers need to do is pull the RJ-45 connector from the 10Base-T concentrator and plug it into the Ethernet switch.

Ethernet switches also can be used as backbone devices, with segments from several departments or on different floors connected to each port. All of the switches except Lannet's support thousands of MAC (media access control) addresses for just this purpose. At present, Lannet only accommodates eight MAC addresses per port — ruling it out for backbone use — but the vendor says an upgrade is on the way.

It's important not to confuse Ethernet switches with port-switching hubs. Port-switching hubs allow net managers to shift users from one segment to another from a central console — without physically reconfiguring the hub.

The Loading Dock

Given that Ethernet switches, by definition, will be used to alleviate saturation problems, it's critical to know how they withstand heavy loads. To evaluate performance, the STbench test suite generates multidirectional traffic between six Ethernet ports on the switches under test. Each port was assigned eight MAC addresses, and frames were randomly switched between ports. Individual tests were run with 64, 128, 1,024 and 1,518-byte frames. All of the devices could handle the two larger frames with no losses; thus, only the test results for the three smallest frames are published here. At each frame size, the test suite generated the maximum number of packets allowed by Ethernet, thus ensuring that the switches were fully loaded.

Three W&G dual-port DA-30s were installed on the testbed to generate and analyse traffic. One analyser and one generator were hooked to each of the six switch ports. All told, the DA-30s recorded a total of 19,000 test results. The production, processing, and analysis of the test data were done in collaboration with a team of engineers from EDF (Electricité de France), the French national power authority.

All of the products were offered a full 10Mbps load, and the percentage of packets lost was determined. The thinking behind this test was that if a switch loses packets it indicates that it can't work at wire speed, a problem on heavily loaded Ethernets.

The test results (see the figure on page 68) are important not just in themselves but for what they reveal about the under-

Ethernet Switches Evaluated (Cont.)

lying technology of Ethernet switches. With the exception of the Lannet and Kalpana products, all of the other devices are actually multiport bridges. They use RISC processors combined with transparent bridging software to switch frames between ports.

The drawback to such an approach is that it makes switching processor-dependent; under heavy loads, performance can drop — at times significantly.

Bridges, by definition, use store-and-forward techniques to pass frames along. When a frame hits a port, the bridging software first stores it and then examines it to determine whether it's to be kept on the current segment or shunted to another.

As part of this examination, the bridge looks at the FCS (frame check sequence) field located at the tail end of each packet. It uses the information contained in the field to perform a CRC (cyclic redundancy check) calculation to see if the packet is good or bad. Alantec, the fastest of the software-based devices — takes 40 microseconds to process each frame and calculate each frame's CRC. That sort of processing overhead must be added to the 51.2 microseconds that Ethernet requires to forward a 64-byte frame onto the media. And 64 bytes is the smallest frame allowed — delays increase with frame size; a 1,518-byte frame, for example, needs 1.21 milliseconds to be retransmitted.

In contrast, the tests reveal that Kalpana's on-the-fly switching — which does not involve store and forward — needs only 40 microseconds to completely switch a packet to its destination. The EPS-1500 simply reads the first part of each frame containing the destination and source addresses and immediately opens a virtual circuit between the appropriate ports. It should be noted that Kalpana's approach prohibits it from performing CRC calculations, which means that if it receives bad packets it will propagate them through the network.

Lannet, in contrast, does use a store-and-forward technique. But the LANswitch also saves time by not doing a CRC calculation. Nevertheless, the product is tuned to pick out misaligned frames (one of the more common types of bad packets), so there's less risk involved.

Unlike the other Ethernet switches, both the Kalpana and the Lannet products are hardwired devices that use ASICs (application specific ICs) to perform the actual switching. The advantage to this scheme is that it ensures that switching is not processor-dependent; under heavy loads, there's no performance fall off.

Like telecommunications gear, the Etherswitch uses a cross-point matrix to establish point-to-point virtual circuits between ports. Lannet employs a high-speed backplane to switch packets among ports.

How Low Can You Go?

The STbench test suite also gauged latency for all the switches. Latency represents the delay incurred by a frame as it passes through a network device. On the wide area, the effects of this sort of delay are negligible when compared with the time it takes to propagate signals over long-distance lines. On the local area, however, latency can have an impact on network performance.

As indicated, the lowest latency for all store-and-forward devices is 51.2 microseconds, the time Ethernet takes to retransmit a 64-byte frame. In other words, even if a bridge could process successive frames instantaneously, it would still delay every frame by 51.2 microseconds.

But even this minimal latency can be a problem when it comes to file transfers. Assume that a workstation sends one acknowledgment frame for every data frame it receives from a server — standard operating procedure with most NetWare applications. The round-trip delay between successive data frames passing through a single bridge would be 102.4 microseconds — effectively reducing Ethernet bandwidth by half.

As shocking as that sounds, bear in mind that it's the best-case scenario for file transfers involving only 64-byte frames; the LANswitch alone of the store-and-forward devices has a latency that equals the Ethernet minimum. The other multiport bridges introduce greater delays, and this can take an even bigger bite out of effective throughput.

True, some higher-level protocols employ windowing techniques that allow 4, 8 or even 16 frames to be sent in a single transmission. By reducing the number of acknowledgments transmitted, these various schemes can reduce the delay between successive data frames. While this alleviates the effectiveness of latency, it cannot eliminate it entirely.

What's more, most file transfers will not consist exclusively of 64-byte frames. In a way, this is actually good news. While it's true that the latency of store-and-forward devices rises as frame length increases, tests indicate that when processor-dependent software switches are fielding longer frames, their latency is within 20% to 30% of their hardware counterparts. That's compared with the 200% to 800% greater laten-

The Test Lineup

VENDOR	PRODUCT	TYPE	MAXIMUM PORTS/ CHASSIS	MAXIMUM MAC ADDRESSES/ CHASSIS	BACKPLANE BANDWIDTH	FDDI	TOKEN RING	FILTERS BAD FRAMES
Alantec Comperex (02) 878 5888	Powerhub	Bridge/router	12	8,192	400Mbps	Yes	No	Yes
Chipcom (02) 416 0653	Galactica	Bridge	32	64,000	1Gbps	Yes	No	Yes
Cabletron Systems (02) 950 5900	ESX-820	Bridge	8	16,000	N/A	Yes	No	Yes
Kalpana NetComm (02) 888 5533	Etherswitch EPS-1500	Switch	15	4,096	N/A	No	No	No
Lannet Toren (03) 242 5050	LANswitch	Switch	128	1,024	1.28Gbps	No	Yes	Some
SMC (02) 238 2206	Elite ES/1	Bridge/router	20	8,192	800Mbps	Yes	Yes	Yes
3Com (02) 959 3020	LinkBuilder 3GH	Bridge	88	11,000	300Mbps	Yes	No	Yes

cy that software switches exhibit when working with small frames. Still, it's almost impossible to accurately predict how these differences will affect end-users on a network since performance depends entirely on the specific mix of long and short frames.

Latency becomes a critical issue when voice and imaging applications are involved, since both are very sensitive to delays introduced by the network components. Multimedia, in general, demands that latency be both low and constant. In addition to having the lowest latencies, Kalpana and Lannet also exhibited the most consistent delays.

When it came to 64-byte frames, the Kalpana Etherswitch exhibited a record low latency of 40 microseconds. What's more, since the Etherswitch processes packets on the fly, latency will increase only marginally with longer packets. Lannet was not far off the record with a latency of 51.2 microseconds for 64-byte frames — equal to Ethernet retransmission time. And it processed 1,518-byte frames in 1.21 milliseconds, once again matching the Ethernet retransmission delay. The 64-byte latency for Linkbuilder was 80 microseconds; the Powerhub came in with a 90-microsecond delay. The Galactica and the ESX-820 both exhibited an average delay of 130 microseconds for 64-byte frames. These test results are only approximate because both switches tended to drop packets, which made it difficult to make exact measurements.

The latency of SMC's Elite ES/1 was way out of line — 770 microseconds for 64-byte frames. It's possible that this level of delay was caused by the ES/1's central processor architecture, which requires all frames to pass through the switch's Packet Processing Engine. SMC has disputed ENL's findings (see below).

HANDLES PORT OVERLOADS	FULL DUPLEX ETHERNET
No	No
No	No
No	No
No	Yes
Yes	Yes
No	No
No	No

Looking for Trouble

Bad frames can be a problem on an Ethernet — even crashing a network if they're allowed to propagate across segments. Runts — frames shorter than the legal 64-byte limit — are typically caused by malfunctioning adaptors. Jabbers — a steady stream of frames longer than the 1,518-byte Ethernet maximum — are usually the result of faulty cabling. In order to see how effectively Ethernet switches deal with these troublemakers, the STbench generates the two most common types of errored packets. It also creates frames with FCS errors.

Only multiport bridges can filter all types of bad frames. And tests indicate that the Elite ES/1, ESX-820, Galactica, Linkbuilder 3GH, and Powerhub do exactly that. Lannet's LANswitch rejects runts and short frames with FCS errors. It also automatically cuts jabbers down to 1,518 bytes. The vendor claims that its product discards misaligned frames, but this capability was not tested.

Packets at the Port

If segments are combined in a switch, there's always the possibility that the aggregate traffic heading for one port will exceed Ethernet's 10Mbps legal limit. When this sort of overload occurs, packets will normally be lost. To see if this is always the case, the STbench generates three levels of traffic, sending 4,000 64-byte frames each second from each of five ports to a single destination port. The total load on the receiving port is thus 20,000 frames per second (fps), well above Ethernet's 14,880-fps limit for 64-byte frames.

Lannet's LANswitch alone of the products tested offers full protection against traffic overheads. In fact, the LANswitch did not drop any frames. The vendor has created a simple but ingenious feature that allows the switch to throttle back the source ports when the destination port is getting too much traffic. The other Ethernet switches simply drop all excess frames when a port is overloaded.

Alantec

The Powerhub combines the functions of a switch, bridge, and multiprotocol router for Ethernet and FDDI. The chassis is limited to 12 10Mbps Ethernet ports, which can be a drawback on very large networks. Very much in its favour, though, is the fact that the Powerhub accepts FDDI modules and can support two FDDI rings. It also routes IP, IPX, and DECnet and translates bi-directional Ethernet/FDDI packets, which makes it particularly versatile.

The Powerhub uses a high-speed shared memory interface between its 12 Ethernet controllers; its RISC microprocessors deliver an aggregate processing power of 45 MIPS (million instructions per second).

Alantec's multiport bridge turned in the best performance of any of the software-based devices evaluated, handling 100% of the 256 and 128-byte packets offered at full wire speed and dropping only 1% of the 64-byte packets. The Powerhub also exhibits fairly low latency and rejects bad frames. When its ports are overloaded, however, it drops frames.

Lannet

The LANswitch was able to handle all packets offered at full wire speed, without dropping a single one. It also exhibited the best latency of any of the store-and-forward devices — 51.2 microseconds for 64-byte frames and 1.21 milliseconds for 1,518-byte frames. Both numbers represent the minimum delay imposed by Ethernet for packet retransmission.

The LANswitch Ethernet modules can be dropped into the vendor's Multinet chassis; the 18-slot LET-36 box is equipped with a 1.28Gbps backplane.

The LANswitch is the only product evaluated that protects against traffic overloads at the destination ports. At present, the LANswitch only accommodates eight MAC addresses per port, although the vendor says an upgrade is imminent. The product also does not filter all bad frames. proprietary full-duplex Ethernet links can be established between chassis.

Standard Microsystems (SMC)

The Elite Switching hub ES/1 is a five-slot chassis that can accommodate Ethernet, Token Ring, FDDI, and HSSI (high-speed serial interface) modules. The ES/1 can route IP and IPX; it also supports source routing and translates between Ethernet, Token Ring, and FDDI. Each Ethernet module features 4 10mbps ports; the chassis accepts five modules (for a total of 20 ports).

The ES/1 exhibited the highest latency of all the switches evaluated — 770 microseconds. The ES/1 finished in the middle of the pack for the performance tests, fielding 97% and 93% of the 256 and 128-byte frames offered. It dropped 10% of the 64-byte frames. The switch rejects bad frames. Gandalf Systems uses the ES/1's technology in its new Prism hub.

Ethernet Switches Evaluated (Cont.)

After SMC had the opportunity to review the test results for its ES/1 it hotly disputed the significance of the latency findings and questioned the test methodology and goals. The vendor indicates that the latency results were 'skewed by an adaptive buffer allocation algorithm . . . designed to minimise frame loss.' According to the vendor: 'Since the ports used for latency testing exhibited lower buffer utilisation rates than the other ports, they became 'donor' ports to the 'needy' flooded ports, increasing the overall latency for the specific ports under test.'

After double-checking the latency results, ENL indicates that the same level of latency was observed on the ES/1 under light and medium loads (4,000 fps and 8,700 fps on the four ports used to stress the switch). SMC says that Version 2.3 of its ES/1 software much reduces latency.

Test Methodology

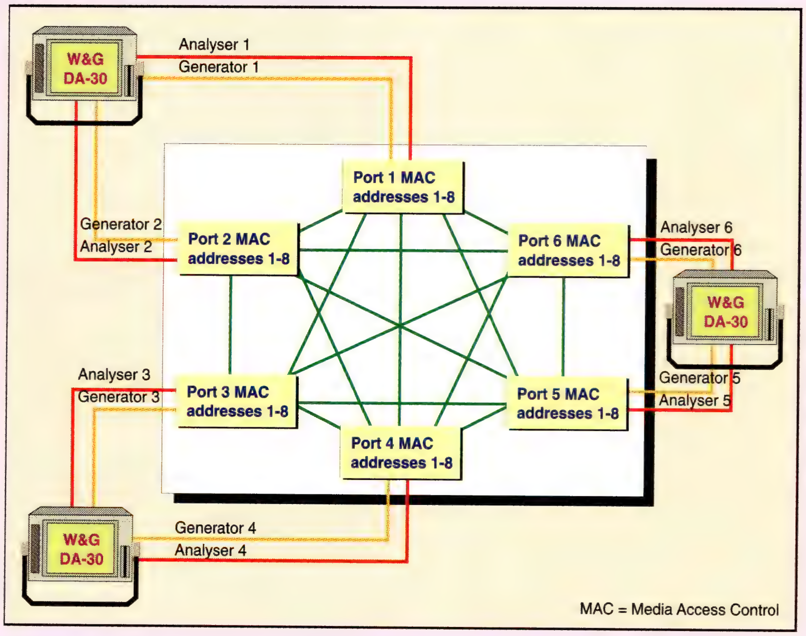

The European Network Laboratories used the most recent version of Switch Test Bench (STbench), developed in cooperation with Jean-Christophe Bestaux of Wandel & Goltermann. It is designed to run on the W&G DA-30 protocol analyser and measure throughput performance for multidirectional traffic and device latency under loaded conditions. STbench runs 18 tests grouped into four classes; five performance tests, four latency tests, six behaviour tests, and three overload tests.

The performance tests measure the rate of packet loss under heavy loads. Individual tests are run for 64, 128, 256, 1,024 and 1,518-byte frames. The tests generate cross-pattern traffic on six Ethernet ports, each of which is assigned eight MAC (media access control) addresses. Source and destination MAC address for each frame are selected at random. Since there are 5 x 8 such destination addresses available to each frame assigned to any of the eight source addresses on each of six ports the total number of virtual circuits established during each test sequence equals 1,290 (5 x 8 x 8 x 6). This burst of multidirectional traffic runs for 30 seconds. Traffic loads were calculated as follows. The maximum rate for 64-byte frames allowed by Ethernet is 14,880 frames per second (fps). Each port is set to send an aggregate of 14,880/2 (7,440) fps to the five other ports. Hence the total traffic sent and received per second by each of the six ports is 14,880 frames. 7,440 frames sent + (5 x (1/5 x 14,880/2) frames received test loads for the four longer frame lengths were calculated in similar fashion.

A new method for measuring latency was developed for this series of tests, designed to subject the device under test to heavy loads. This technique permits latency measurements to be made with 32-microsecond precision. The latency test was run on six switched Ethernet ports using 64- byte and 1,518-byte frames. Blocks of six contiguous frames are earmarked, time-stamped, and then sent from port A to port B at 4000, 8700 and 14,000 fps. Ports A and B are isolated from the four other ports. In order to keep the buffers of the devices under test filled, every block of six time-stamped frames is randomly mixed with 100 filler frames. The test is set to last three seconds. At the same time, cross-pattern traffic is generated between ports C, D, E, and F at the full media rate using 64-byte frames.

STbench also examines the behaviour of Ethernet switches when bad frames circulate on a network. The protocol generates 40, 64, 1,518, and 3,000-byte frames containing errors in their FCS (frame check sequence) fields. Runts (frames shorter than the legal 64-byte limit) and jabbers (a stream of frames longer than the 1518-byte Ethernet maximum) are both generated. The overload tests analyse switch behaviour when two or more ports send frames to a single receiving port at a rate greater than the Ethernet 10Mbps maximum. This test generates three levels of unidirectional traffic between five sending ports and one receiving port at 5 x 2,000, 5 x 3,000 (slight overload), and 5 x 4,000 (heavy overload) 64-byte frames per second.

Chipcom

The Galactica Super-Switching Hub is a four-slot chassis that can accept four 8-port Ethernet switching modules (for a total of 32 ports) or three Ethernet modules (24 ports) and an FDDI module. The Ethernet modules are actually multiport bridges that conform to the IEEE 802.1 standard. Each module employs a 50-MIPS RISC processor to handle frames. It incorporates a 1Gbps high-speed backplane. Frames whose source and destination addresses are on ports of the same module do not circulate on the backplane, which helps keep traffic from backing up.

The Galactica performed acceptably with 256 and 128-byte packets, but dropped 21% of the 64-byte frames. Frame loss also made it difficult to accurately establish the product's latency, but the 130-microsecond approximation puts it in the middle of the pack.

Chipcom OEMs the Galactica from US vendor Artel. At the time of going to press a merger was pending between the two companies.

Cabletron

The ESX-820 is based on the same technology used in Artel's Starbridge Turbo Ethernet Switch, which also is OEM'd by Xyplex. The Cabletron Ethernet modules are OEM versions of those used in Artel's Galactica. The chassis supports eight Ethernet ports, ruling it out for large networks. Slide-in AUI and RJ-45 connectors help simplify installation.

Of all the products tested, the ESX-820 exhibited the worst frame loss, dropping 16% of 128-byte packets and 25% of 64-byte packets. Because the ESX-820 dropped so many frames, it was impossible to gauge latency with complete accuracy. The estimate of 130 microseconds with 64-byte frames puts the ESX-820 in the middle of the pack.

Several months ago, Cabletron introduced a new Ethernet switching module for its MMAC hub. Six of the ESX-MIM's seven ports can be dedicated to Ethernet; the final port supports FDDI or Token Ring (ATM is in the works). The MMAC

DOS

MACINTOSH

UNLEASH THE FULL POTENTIAL OF NETWARE SAA.

If you have Novell® NetWare for SAA™, you know you have a truly powerful LAN-to-host solution. But you aren't taking full advantage of all that power without equally powerful client software.

Unlike basic 3270 emulators, IRMA™ WorkStation for Windows offers features that bring a whole new meaning to functionality and ease-of-use in micro-to-host communications solutions. It's a powerful application development tool, supporting HLLAPI and APPC and including cross-platform API Tool Kits like our own QuickApp™. Its productivity tools make virtually every task easier and faster.

And since IRMA WorkStation for Windows complies with DCA's Universal Communications Architecture, a framework for consistency across hosts, transports and operating systems, your investment will be protected as your connectivity needs change.

IWW supports many other gateways and connections, including TN3270 over TCP/IP. But if you only need NetWare for SAA support, we have IRMA NetWare WorkStation for Windows.

It's optimized for NetWare for SAA, offering IWW's functionality at a lower price.

And now you can get both the NetWare for SAA gateway and IRMA WorkStation for Windows, DOS, OS/2, Mac and Windows NT™ from DCA®.

So you have a single source for sales, training and support. For a limited time, you can trade up from NetWare 3270 LAN WorkStation™ or any other micro-to-host software for only $180. (That's a savings of up to $400. And MultiPaks save you even more.)

The combination of NetWare for SAA and IRMA WorkStation for Windows can bring an easy-to-use graphical interface to multiple mainframe applications.

We'll also send you a free demo disk of any IRMA WorkStation software for any platform.

If you want the best LAN-to-host solution, then IRMA WorkStation for Windows is for you.

Call DCA on (02) 451 8255 or fax the coupon below for your $180 trade-up or free demo disks.

Only QuickApp allows developers to create GUI interfaces for legacy applications without knowing HLLAPI.

The only Windows-to-host software with a QuickBar for instant access to functions like File Transfer and Print, as well as other 3270 commands like Reset, Clear and Attention.

FAX TO DCA ON (02) 451 4034 TODAY.

PLEASE TICK ☐ *Yes!* Please send me my free demo disks for IRMA WorkStation.

PLEASE TICK ☐ *Yes!* I would like to upgrade to the IRMA WorkStation for $180.

I have the following computer platform: ☐ Windows ☐ DOS ☐ Macintosh ☐ OS/2

NAME ____________ TITLE ____________

COMPANY ____________ TEL ____________

ADDRESS ____________

POSTCODE ____________

OS/2

WINDOWS

PC MAGAZINE EDITORS' CHOICE February 26, 1991 IRMA WorkStation for Windows, Ver. 1.0

yes It runs with NetWare

COMPAQ Qualified Solution

DCA

FREEDOM TO COMMUNICATE

DCA Australia Pty Ltd, A.C.N. 056 817 086 20 Rodborough Road, Frenchs Forest, NSW 2086.

Or call (404) 475 8380. Trade-up good in U.S. and Canada only to 12/31/93. Allow 6–8 weeks for delivery. ©1993 Digital Communications Associates, Inc. All rights reserved. DCA is a registered trademark and IRMA and QuickApp are trademarks of Digital Communications Associates, Inc. Windows and Windows NT are trademarks of Microsoft Corporation. The Compaq Qualified Solution Logo is a trademark of the Compaq Computer Corporation. All other trademarks are the property of their owners. Developer-tested only. Novell makes no warranty with respect to this product.

MATRIX D774

Ethernet Switches Evaluated (Cont.)

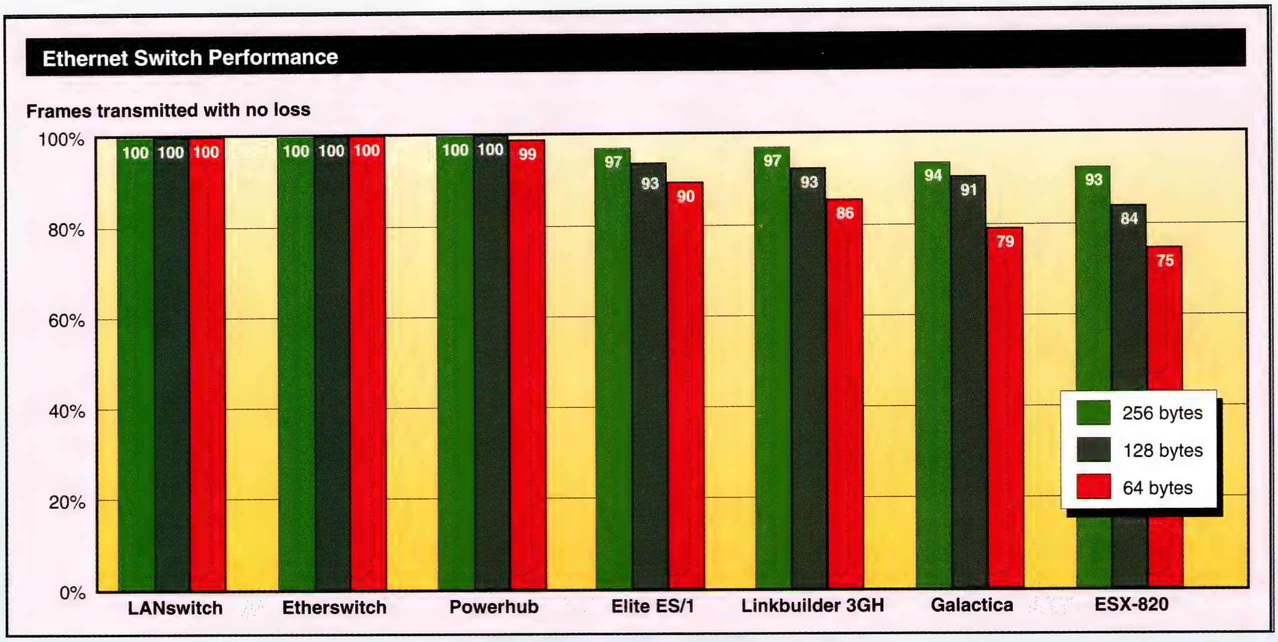

chassis can be had with three, five, or eight slots; backplane bandwidth is 630Mbps.

Kalpana

The Etherswitch EPS 1500 tied for first place in the performance tests and exhibited the lowest latency of any product evaluated — 40 microseconds. Kalpana's impressive performance can be traced directly to its architecture: the Etherswitch features a cross-point matrix switch implemented as an ASIC. What's more, this dedicated hardware component reads only the first 12 bytes of each Ethernet frame before setting up a virtual circuit and starting to transmit between the appropriate source and destination addresses. By doing so, the Etherswitch can actually shave more than 11 microseconds from the 51.2-microsecond retransmission delay imposed by Ethernet on 64-byte frames. Because the Etherswitch processes packets on the fly, however, it does not perform CRC calculations and consequently can pass bad frames onto the network. The switch also drops frames when the ports are overloaded.

The Etherswitch is a 15-slot chassis. If more ports are needed, units must be chained. This can be done via 10Mbps half-duplex connections; the Etherswitch also allows 20Mbps links via its proprietary full-duplex Ethernet implementation. Each port supports up to 170 MAC addresses, which makes the Etherswitch suitable to backbone networks.

Two major networking vendors, Hewlett-Packard and Synoptics Communications, both OEM the Kalpana switch.

3Com

Linkbuilder 3GH is based on technology used in the LANplex switch from Synernetics. The Linkbuilder's 300Mbps backplane actually consists of three FDDI rings. The 11-slot chassis accommodates two types of FDDI modules as well as 8-port switched Ethernet modules (for a total of 88 dedicated Ethernet ports). A translating bridge links the Ethernet and FDDI modules.

The Linkbuilder performed acceptably with 256 and 128-byte frames but lost 14% of the 64-byte frames offered. At 80 microseconds, it had the second-lowest latency of any of the store-and-forward devices and only twice that of the record-setting Etherswitch). The Linkbuilder filters bad frames. It can learn up to 8,000 MAC addresses, which helps suit it to backbone networks.

3Com also had several comments once it had the chance to review the results for its Linkbuilder 3GH. First the vendor does not believe that 'the 64-byte packet test represents the characteristics of a real network,' adding that 'rarely do many 64-byte packets fly through a network one after another.' The vendor believes that 'a real network' would have 'variable packet sizes that average between 300 and 1,024 bytes.' Second, 3Com states that: 'The 3GH's latency between packets (not including packet storage time) has been measured at 25 microseconds.' Finally, the vendor believes that 'features such as IP fragmentation, CRC checking, broadcast storm firewalls, and filtering' should be included in product evaluations.

When tested with larger frames (1,024 and 1,518 bytes) the Linkbuilder did, indeed, perform at wire speed with no packet loss — as did all the other switches evaluated. Steady streams of 64-byte packets are not generated by STbench; rather, it creates random patterns of multidirectional traffic to ensure that every port is tested equally for sending and receiving frames. Further, in audits of four of five major production networks, ENL has observed that small frames represent more than 70% of the traffic.

ENL's findings would tend to confirm 3Com's claim of 25 microsecond latency if packet storage time is not included in the measurement. When 51.2 microseconds is allowed for packet storage, 3Com's latency for 64-byte frames would be 76.2 microseconds (25 + 51.2 microseconds) — not far off the 80 microseconds that ENL measured. These measurements were taken under loaded condition, which may account for the discrepancy.

Finally, features like firewalls and filtering are found on bridges. Since this was a switch evaluation, they were not considered. In fact, tech support representatives from the multiport bridge vendors went to great lengths to disable such functions, believing that they might slow switch performance.

Robert Mandeville is Director of European Network Laboratories (ENL), a vendor-independent organisation that tests network hardware and software. This product evaluation was carried out in cooperation with Electricité de France, the French national power authority.

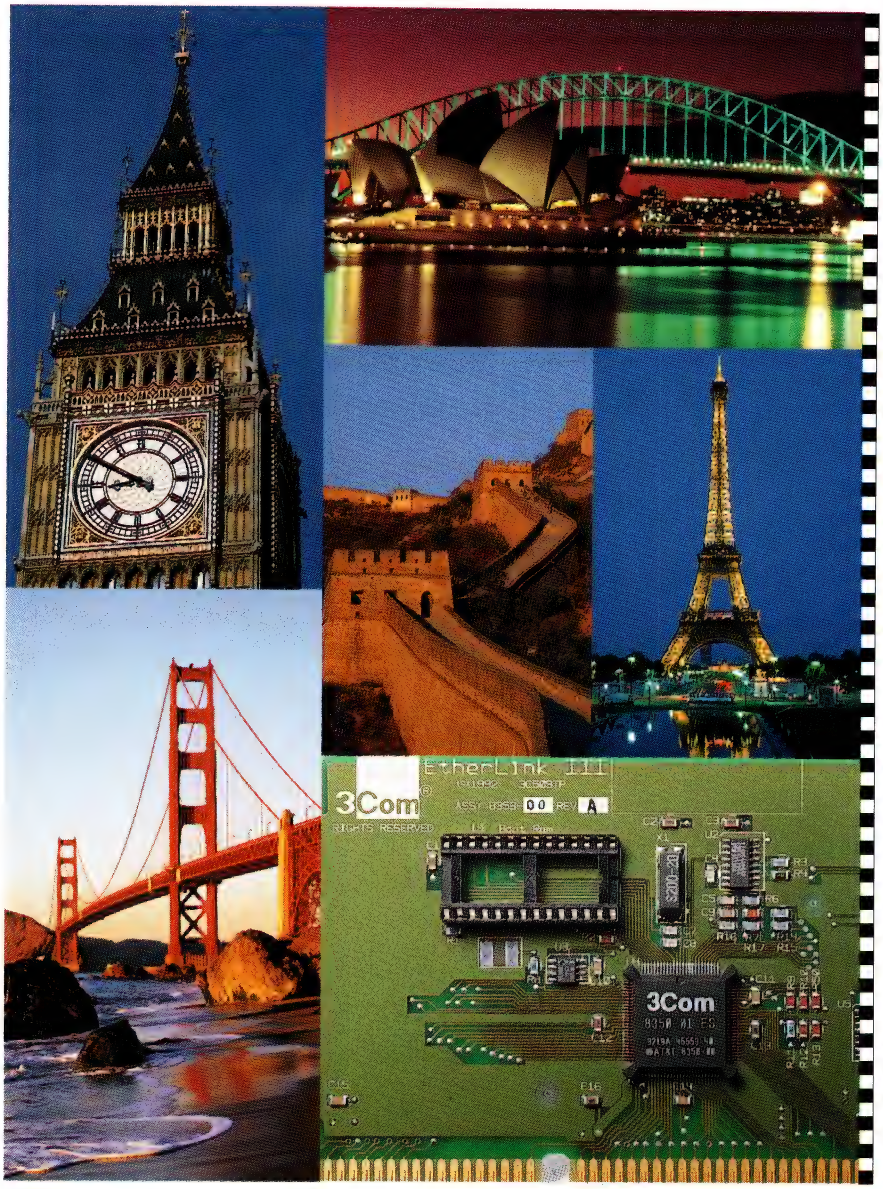

There's only one original Parallel Tasking.™

■ What would you call a technology that delivers the fastest adapters in the world? That has won immediate acceptance from network managers around the world?

Millions know it by the name Parallel Tasking. And you'll find it in only one place: 3Com.®

With 3Com's Parallel Tasking adapters there is simply no faster way to run your network. And there is no other adapter line that offers more convenience and more performance at a lower price.

An EtherLink® III Parallel Tasking adapter is SNMP-manageable. It is auto-configuring. And every 3Com EtherLink III adapter is backed by a lifetime warranty.

It's no surprise, then, that more than two million have been sold.

But why not judge for yourself? Call **(02) 959-3020, ext. 20** for your free *Guide to Benchmarking Network Adapters* from LANQuest Labs, and complete information on 3Com's EtherLink III family of adapters.

It's never been easier to put unparalleled performance on your network.

3Com®

Networks That Go the Distance™

©1994 3Com Corp. (408) 764-5000. 3Com, Networks That Go the Distance, EtherLink, and Parallel Tasking are trademarks of 3Com Corporation. LANQuest Labs is a trademark of LANQuest Group. All other trademarks are property of their respective owners.

from page 62

packets per second, which means that any one packet can be removed from any port's input buffer to its destination port's output buffer in about 12 microseconds.

This means that reception of the preamble for the minimum size frame (64 Octets inclusive of preamble and SFD) will be underway. Each port's receive buffer is cleared before reception of the incoming packet's preamble and SFD is complete — by just over 2.5 microseconds, some of which is taken up doing a CRC check and an address table look-up — and transmission of the 1,518 byte packet can begin immediately (9.6 microseconds) after reception of the incoming 64 Octet packet ends.

This means that, in so much as pure packet forwarding goes, the Galactica has latency (i.e. propagation delay between ports) figures of effectively one frame time under any conditions. The Galactica's 1Gbps bandwidth is the minimum necessary to guarantee this. Galactica can also be configured to act as an FDDI packet switch, but with only four ports.

Digital's Gigaswitch, for which similar calculations can be made, is designed to provide this kind of performance for up to 22 FDDI ports and has a 3.6Gbps backplane. The Gigaswitch has a forwarding rate of 290,000 FDDI packets per port per second and a filtering rate of 446,000 FDDI packets per port per second. These devices are built to achieve this level of performance specifically to keep their latency figures within very narrow limits. Remaining within these limits is necessary to avoid a negative impact on throughput.

This is why administrators who have deployed slower Ethernet packet switches sometimes see an increase in collisions and dropped packets on their networks and consequently a much lower network through-

CSMA/CD and All That

Carrier Sense, Multiple Access with Collision Detection is arguably what defines Ethernet. The simplicity of this mechanism is one reason why Ethernet is so cost effective, while its consequences are the source of many of Ethernet's physical limits. While CSMA/CD has been the root cause of many an administrator's frustration the fact that it remains basically unchanged from its original specification in 1975 is a testament to its overall reliability and robustness.

The basic state of any station in a CSMA/CD environment is simply monitoring the transmission channel for the presence of a *carrier* signal. This state is referred to as *carrier sensing* and is physically the circuitry of a transceiver monitoring the level of electrical (or optical in the case of Fibre Optic Inter-Repeater Link (FOIRL) or 10Base-F) signals in the network cable to which it is attached. When a carrier is detected a *CarrierSense* signal is asserted and the station defers any of its own transmissions. When the CarrierSense signal is no longer asserted the station continues to defer its transmissions for a period of time referred to as the *interpacket delay* and then either returns to a carrier sensing state or, if it has a packet waiting to be transmitted, begins transmission irrespective of the CarrierSense signal. If more than one station has a frame awaiting transmission at the end of the interpacket delay all of them will attempt to transmit at once (a *multiple access*) and a collision will result.

Collision detection is achieved in much the same way as CarrierSense and a *CollisionDetect* signal is asserted. All stations at which a CollisionDetect signal has been asserted then *enforce* the collision by transmitting any sequence of bits (except a CRC on any transmitted packet fragment) of between 32 and 48 bits. This bit sequence is referred to as a *jam* signal and is intended to ensure that all stations on the network notice the collision. After jamming the network all stations which have detected a collision *backoff* for a random number of slot times (the slot time for Ethernet is defined as 512 bit times) before attempting re-transmission. Retransmission will be attempted up to 15 times with the range from which the random delay is selected increasing for each attempt (this process is called Truncated Binary Exponential Backoff) before the attempt is abandoned on the assumption that the channel has failed or is completely overloaded. The range, r, from which an integer is selected at random for the nth retry is given by $0 <= r <= 2^k$ where $k = \min(n,10)$.

Unfortunately the pristine simplicity of all this is necessarily soiled by reality. Firstly signals in a cable — copper, or optical fibre — take a finite amount of time to travel between two points. In the case of the various copper cables used for Ethernet in a coaxial system this propagation speed is anywhere from 65% to 77% of the speed of light in vacuum — about 300,000km/second. It also takes time for a signal to reach a level at which a transceiver can detect it and it takes time for the electronics in that transceiver to make a transition from one state to another — say from CarrierSense = False to CarrierSense = True — or having detected a signal to synchronise themselves with the framing (referred to as *Acquisition* of the signal) and thus actually detect incoming bits.

These later delays are the reason for the interpacket delay (not less than 9.6 microseconds and not more than 10.6 microseconds) and the transmission of preamble (7 Octets) and a Start of Frame Delimiter (SFD 1 Octet). There may also be repeaters, to which these delays also apply, in any given signal path in the network. When all these delays are taken into account the worst case *round trip delay* — i.e. the time a signal from one station will take to reach another station, be detected and return to the its source — is just over 46 microseconds.

To eliminate the possibility of a station transmitting a complete frame before all stations can detect a collision the minimum packet size must be such that it takes longer than this round trip delay time to transmit. It follows that all stations must be capable of signal acquisition within the time required to transmit this minimum size packet and that frame fragments generated from collisions must be less than minimum packet size. These considerations, and the fact that the backoff delay in case of a collision must be an integer multiple of a time greater than the worst case round trip delay, constrain the minimum packet size on an Ethernet to be more than 464 bit times (the time required to transmit 464 bits). The minimum packet size on an Ethernet has been defined as one slot time which, as mentioned above, is defined as 512 bit times.

Note particularly that throughput in a CSMA/CD environment is dependent upon the number of stations which attempt to transmit at once and the number of times which, on average, each station is required to retry a transmission. The size of a packet, and thus the time required to transmit it, simply influences the probability of there being a single collision directly after its transmission. Repeated collisions, which are what cause an Ethernet's response time to rise exponentially, can only occur as the result of there being sufficient active stations on a network to ensure that there is a high probability of two stations attempting a simultaneous transmission.

Once this happens a relatively small increase in the number of stations may result in a higher number of retries after collisions. This in turn increases the average transmission delay and results, generally long before CSMA/CD's 15 retry limit is reached, in protocols and applications timing out. Administrator's Moral: keep a close eye on your network's peak usage figures and collision rate.

Graeme Le Roux

IF YOU THINK FITTING A FULL-FEATURED NETWORK ANALYZER INTO A NOTEBOOK IS AMAZING . . .

. . . HAVE A LOOK AT THE NEW EXPERT SNIFFER

You supply the notebook and let us supply our powerful award-winning technology. Available for both Ethernet and Token Ring, it uses PCMCIA-based software. This lightweight yet sophisticated analyzer will enable identification of network problems, pinpointing their origins and automatically suggesting solutions. The Notebook Sniffer Analyzer runs on a variety of platforms at a price to suit your budget.

CON TECH

INTERconnect

For more information call -

Australia - Network General Australia Pty Ltd
(02) 911 7770
New Zealand - INTERconnect
(09) 358 0322

Network General

THE DATA AND TELECOMMUNICATIONS MANAGEMENT MAGAZINE

Why not get the WHOLE story on networking?

Whether it's FDDI, OSI, ATM, ISDN, SNA, Frame Relay, Fastpac, Token Ring or 10Base-T, you need to be on top of the rapid changes now taking place in the communications industry.

In the pages of Australian Communications, every month the region's top communications journalists, leading industry figures, academics, professionals and industry experts bring you the latest news and analysis on the technical, political and management issues facing this dynamic industry.

The time to subscribe is NOW! with savings up to 20% off the newsstand price.

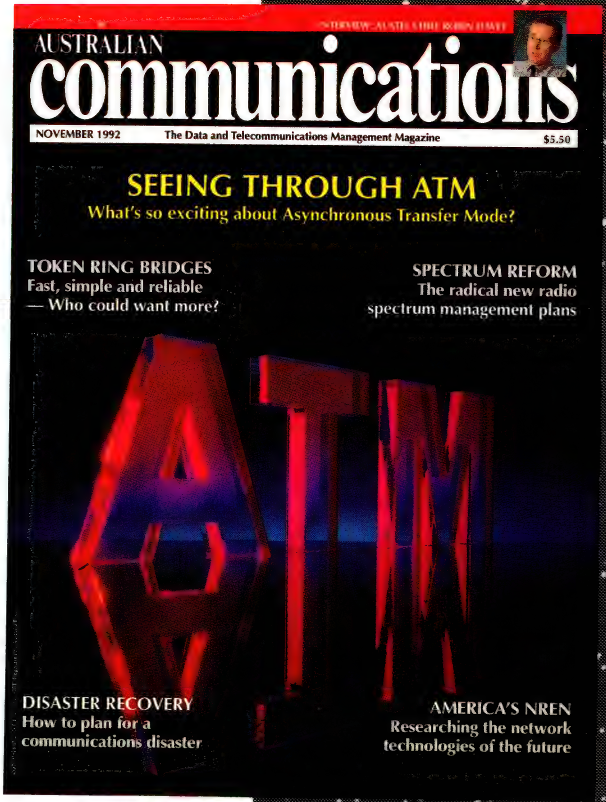

AUSTRALIAN
communications
NOVEMBER 1992 The Data and Telecommunications Management Magazine $5.50

SEEING THROUGH ATM
What's so exciting about Asynchronous Transfer Mode?

TOKEN RING BRIDGES
Fast, simple and reliable
— Who could want more?

SPECTRUM REFORM
The radical new radio spectrum management plans

DISASTER RECOVERY
How to plan for a communications disaster

AMERICA'S NREN
Researching the network technologies of the future

SUBSCRIBE NOW

☐ 22 issues, two years, for $96 (save 20%)
☐ 11 issues, one year, for $54 (save 10%)

I enclose my cheque for AUS$..........

☐ Bill my company Authorising signature..........

Charge my: ☐ Amex ☐ Diner's ☐ Visa ☐ Bankcard ☐ Mastercard

No .... Expires..........

Name ....

Company ....

Address ....

.... Postcode ....

Signature ....

Post to: Australian Communications, Level 4, 541 Kent Street, Sydney NSW 2000
OR: Fax your application to us on (02) 264 2244

put than they might have expected after deploying the new devices.

Latency figures aside, the Gigaswitch, Galactica, and all other packet switches, are logically an ISO/OSI bridge — albeit one with several parallel logical paths at any one time. They are protocol independent devices which make forwarding and filtering decisions on the basis of MAC layer addresses. MAC layer addresses are stored in an address table (common to all ports) and may be aged on the basis of metrics such as usage and time. Some Ethernet packet switches, including the Galactica, will permit you to set custom filters, although this will increase latency and static addresses. Many also support the Spanning Tree Algorithm (STA) just as most ordinary bridges do. One minor trade-off with all types of switched network is that the Spanning Tree Algorithm can take significantly longer to converge, simply because there are many more paths, and therefore potential loops, within a switched network for the STA to check.

Not Without Limits

While Ethernet packet switches can significantly increase the throughput of a network they have their limitations. One of the assumptions made by designers of Ethernet packet switches is at that, at any given moment, most of the packets being received by the unit will be destined for different destination ports. This means a minimum number of broadcasts which have to be forwarded to all ports except the one on which they are received and therefore cause the Ethernet packet switch to act as a repeater/concentrator.

This assumption is also invalidated in an environment where most packets received have the same destination — for example a PC server. Unfortunately most PC networks have a significant proportion of broadcast traffic, much of it between many workstations — stations acting exclusively as clients — and few servers. Even if you had a network full of workstations acting as described in the example above, the cost difference between stackable hubs (about $100 per port these days) and the Galactica as configured above ($2,848 per port with dual power supply and software — remember you are buying the equivalent of 32 bridges) would put the company beancounter into low orbit.

What you can do with an Ethernet packet switch is connect your servers and stackable hubs to it, as shown in Figure 3. If your servers have only one interface each they can only serve one workstation at once so, assuming equal access loads, the best you can do is make provision for as many paths through your switch as you have servers.

What will really hurt the throughput of the network shown in Figure 3 is any operation which ties up a server's network interface for a long period of time; for example, a server-to-server back-up or one server sending a long print job to a printer attached to one of the hubs. An obvious solution would be to set-up a separate network for printing and another for server to server backups.

Figure 4 illustrates such a network, three *private* (or dedicated) networks have been created — one for each printer and one for server-to-server back-up.

A private network is one in which a single Ethernet segment is dedicated to just two stations, therefore a workstation, for example, has a guaranteed 10Mbps channel to a server. Private Ethernets do not guarantee 10Mbps throughput, they simply guarantee the potential for the maximum possible throughput of the slowest station of the pair attached to them. Note that in the configuration shown, the server's software must not provide bridging functions between network interfaces and, depending upon the protocols in use, server routing may need to be disabled and custom filters may be required in the Ethernet packet switch.

In an Ethernet environment a private network is the best you will be able to achieve for streaming applications such as backups. Private Ethernet will not be as efficient as Token Ring or other synchronous access schemes such as FDDI for such applications, but it won't cost as much either. The main problem with private Ethernet is that it becomes impractical when, for example, you need to back-up more than one or two servers.

Switching Types

Since most larger sites have at least two tape units and at least three or four servers in total, what is required is some way of configuring a private Ethernet according to which server you intend to back-up. This is known as *segment switching*. Segment switching allows a network manager to physically remove an Ethernet segment and all the stations on it from one network to another at will — rather like a patch panel you can configure remotely via a management agent. A segment switching hub achieves this by permitting an administrator to logically connect a single port on a single module to another port, or ports, on the same module. These switched segments do not pass data across the backplane of the hub in which the module to which they are attached resides.

If a port is logically connected to a port, or ports, on another module within the hub or across a private network (between two hubs; no bridges, routers or strictly speaking repeaters) to a port(s) on another hub then the term used is *port switching* to underline the distinction — although the terms are sometimes used interchangeably.

Port switching describes a one-to-one-or-more mapping. Many-to-one-or-more mappings are referred to as *module* (all the ports on a single hub module) and 'bank' (a subset of the ports on a single hub module) switching. Figures 5a through 5d on page 77 illustrate segment, port, module and bank switching. Some examples of hubs which are capable of segment, port, module and/or bank switching are the Digital's DEChub 900 Multiswitch, 3Com's LinkBuilder MSH, the SynOptics Lattis System 5000 and the Chipcom ONLine and ONCore families.

Virtually Solved

Where packet switching is used to provide a number of concurrent data paths within a network and therefore an increase in available bandwidth, other forms of switching

BANDWIDTH ON DEMAND
Networking by Dataplex
For networking solutions that work, call us on
(03) 210 3333 (02) 349 3911 (06) 285 3657
Local, Wide Area and Switched Networks
Dataplex
DPX 185 B

Figure 2: Collapsed Backbone

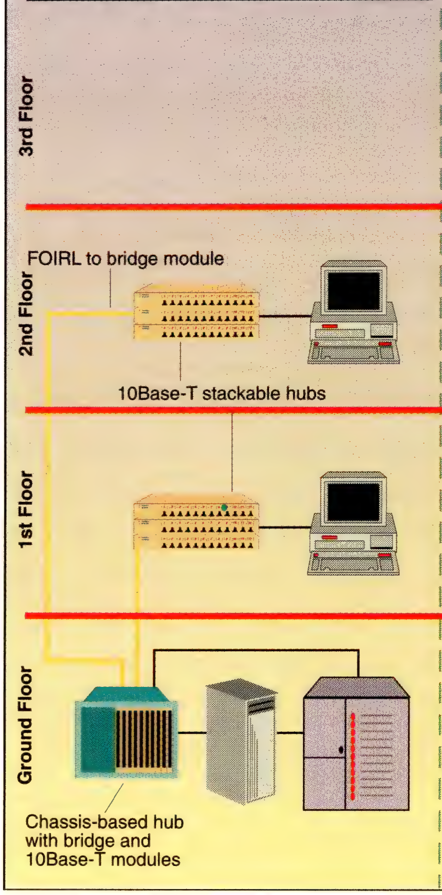

Figure 3: Packet Switched Ethernet

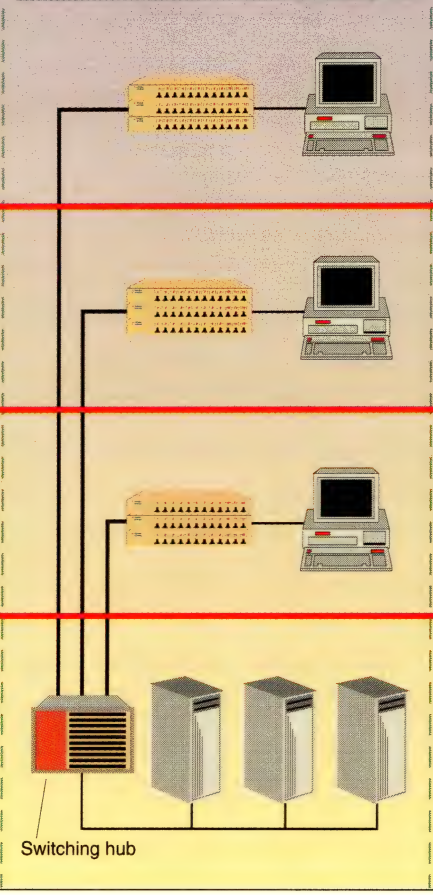

Figure 4: Private & Switched Ethernets

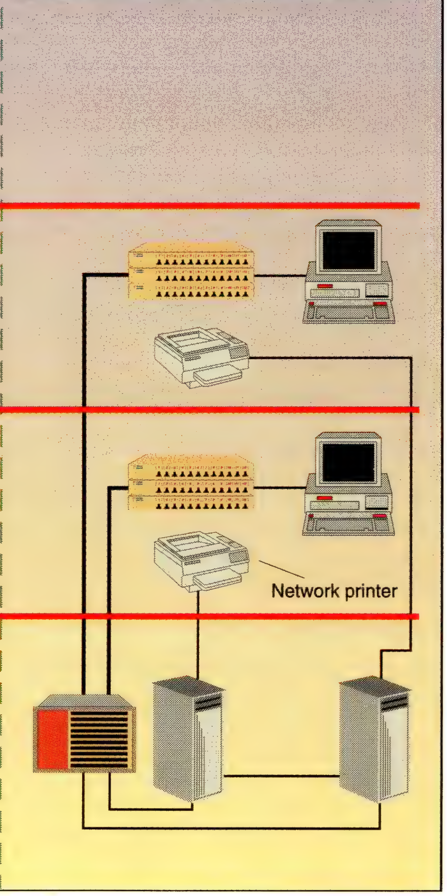

enable administrators to deploy that bandwidth more efficiently. Port and segment switching are best used to provide maximum possible bandwidth for operations like server-to-server file replication, back-up or software upgrades without affecting an operational network's throughput, or for isolating certain groups of users. Module and bank switching are best for coping with moves and changes, setting up security zones and *virtual workgroups*. The virtual workgroup example on page 76 demonstrates how these various types of switching interrelate. The concept of the virtual workgroup can be extended to create a *virtual network*.

While the terms virtual workgroup and virtual network are often used interchangeably, the difference between them becomes apparent, particularly in an IP environment, in the way in which the virtual workgroup or network is connected to the rest of an organisation's network.

A virtual workgroup is generally connected to a network via a bridge, and therefore has the same network address as the network to which it is connected, while a virtual network is connected to the rest of the system via a router — requiring different network addresses (or sub-nets).

Virtual networks are particularly handy in dealing with situations where a group, or part of a group, of users physically move — say to another floor of a building or another building at the same campus. A virtual network can be logically moved to another module in a hub or another hub without having to alter the IP addresses on each user's PC and all the printers or physically move servers. This 'logical' move can be done remotely and, if an administrator chooses, automatically at some scheduled time such as midnight on Sunday. Virtual networks and workgroups require careful forethought as it is unwise to move some types of IP addresses — like management agents.

Any device which is manageable via SNMP has an IP address, part of which is a network address. If you move that address to the other side of a router your management software may 'lose' the device from the network. Most management software would interpret such an event as a network fault and act accordingly. Users on your virtual network would, in most cases, not notice a problem; however your operations staff would have an opportunity to test their troubleshooting technique.

Your users would not be so lucky in the event that you moved a virtual network or workgroup in a secure environment. Ethernet security schemes such as port jamming, as used by the DEChubs, or data replacement, as used by 3Com's LSA, function by mapping a given MAC layer address to a physical port — say on a 10Base-T repeater. A packet destined for a given address is sent to the port to which its MAC layer address is mapped while all other ports see an Ethernet jam signal (port jamming) or a packet containing rubbish (data replacement). This is known as maintaining network privacy.

Unless your logical move re-maps address-port pairs your users will not receive packets at their new ports. It is also possible, should any user attempt to transmit at their new locations, that the security software will interpret this as an attempted intrusion, shutdown the user's port and notify the administrator. It follows from Murphy's Laws of Networking that the administrator will not get this notification or be able to fix the problem when users call because the hub's SNMP agent will be 'lost' from the network map as discussed above.

There are also interesting variations on this theme such as the ability to assign multiple IP addresses to a given devices, while necessarily maintaining just one MAC layer address. This allows, for example in a Chipcom hub, an administrator to Telnet to a Management Module on a given network at a given IP address, change the module's active network and lose the Telnet session — the management module will appear on the newly assigned network with a new IP address. This has to happen

Uncovering network problems one segment at a time could leave you open to some unpleasant surprises.

HP NetMetrix distributed monitoring solutions: Now nothing can hide from your view.

Here you go again. There's a hot spot in your network and you don't know what's causing it. So you begin searching. Segment, by segment, by segment. Meanwhile, the source of your problem continues to raise havoc.

What can you do about it?

Put an end to your search with HP NetMetrix. Our distributed monitoring and analysis solutions allow you, from a single computer console, to instantly know what every segment of your network is doing – whether it's across the hall or across the world.

Take our correlated viewing capability. It shows you how activities on different segments affect each other. So you can quickly find problems based on insight, not intuition.

Of course, your goal is to catch problems before they become a disaster. With HP NetMetrix, you can set networkwide alarms to alert you at the first sign of trouble.

Flexibility and compatibility are other HP advantages. All our solutions are open and are based on the SNMP and RMON standards. They integrate with all leading network management platforms, such as HP OpenView and SunNet Manager, and support a variety of computer hardware platforms.

Best of all, because having a complete view helps you better understand your network, you can make more informed architectural decisions and optimise your network's performance.

Want proof about how HP NetMetrix can help you uncover problems and manage your network like never before? Call our Customer Information Centre on **13 1347** (Australia wide) and ask for extension **2902** for literature on our distributed monitoring solutions.

You'll never look at your network the same again.

A better way.

HEWLETT®
PACKARD

An Example Of Virtual Workgroups

A common administrative headache arises when a department, say accounts, requires a security wall between itself and the rest of the company. All of its hardware has to be physically locked up in a secure area which cannot be reached from the rest of the company's network. Of course this department still wants the network group to administer 'their' network and provide back-up services and access to other users via e-mail.

This is simple enough; one possible solution, using Windows (NetWare would work just as well), is shown in Figure A. Workstations are configured with DOS 6.2, Windows for Workgroups and run only the non-routable NetBEUI protocol. Workstations and the department's Windows NT Advanced Server(s) are connected using 10Base-T as a virtual network using bank and/or module switching within a single hub as LAN #1. LAN #2 is connected via bank switching and consists of a number of printers with 10Base-T interfaces and the department's server(s). LAN #2 runs only the non-routable DLC protocol. LAN #1 and #2 might be combined, however DLC creates reasonably large numbers of frames and can degrade network performance when used heavily. Note that several of these printers might be combined as a printer 'pool' — i.e. a single logical device as far as users are concerned — to provide fault tolerance and a higher system throughput.

Figure A: A Secure Department

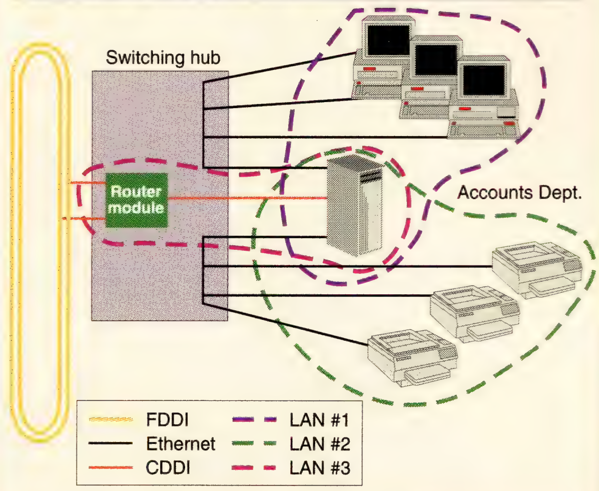

Figure B: Virtual Workgroups/Networks

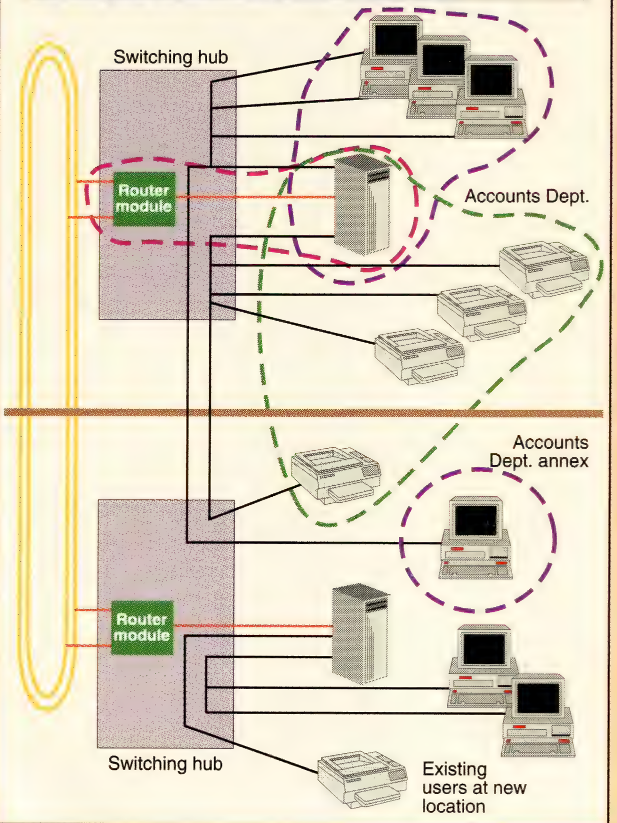

LAN #3 is a private Ethernet running only TCP/IP which consists of the department's server(s) and a router which is internal to the hub and provides a connection to the rest of the company's internet for e-mail, administration and back-up. While LAN #3 could be an Ethernet, FDDI — or more specifically CDDI from server to the hub, and the hub connected to an FDDI backbone — as shown would be a better technical choice, since it is a synchronous access scheme. Notice that the server is physically multi-homed; either three 10Base-T interfaces or two 10Base-T interfaces and a single F/CDDI interface are installed. The server(s) provide e-mail services, data storage, a platform for client server applications, etc. for the department.

In a NetWare environment the server's bridge/routing functions would be disabled in this media configuration. Since Microsoft did not put this basic feature into NT Advanced Server 3.1, the issue does not arise. Departmental traffic is contained and there is no access from outside the departmental domain other than via management stations running TCP/IP and subject to the departmental domain controller's security scheme. If the rest of the company follows this type of topography and protocol set-up, security is consistent, reasonably simple and reasonably high without impeding users. So what happens when the department grows and must physically move some users to another floor?

Without a switched network this would require a virtual duplication of the original network at the new location and a potentially complex rearrangement of route tables. Further, the site's administrators would undoubtedly hope that nobody in the new location would ever attempt to share local workstation resources via Windows for Workgroups, as this would require a complex reconfiguration of workstations to allow them to use TCP/IP and NetBEUI or a bridged connection to the new location.

In a switched Ethernet environment the problem is much easier to handle. The existing hub at the new location is re-configured such that a bank of its available ports are connected to the bank/module supporting LAN #1 at the original location and another port/bank at the new location is similarly switched to LAN #2 at the old location to support printing. Appropriate configuration changes are made to the user's server accounts and the new location is up and running without any effect on other users — either in or out of the accounts department.

The only topographical change to be made — aside from moving PCs — is to run two new backbone cables (possibly just 10-Base-T) between the existing hubs. An alternative would be to set-up such backbones by using tunnelling if the feature was supported by the FDDI router modules — effectively setting up two extra channels with fixed end points over the FDDI ring. But note that if this is done on a regular and *ad hoc* basis the network can degenerate to a complex mesh topography. The beauty of ATM when it matures will be the ability to divide a physical cable into discrete channels of different bandwidths carrying different logical networks — in this case two Ethernet and one FDDI.

Graeme Le Roux

Figure 5a: Segment Switching

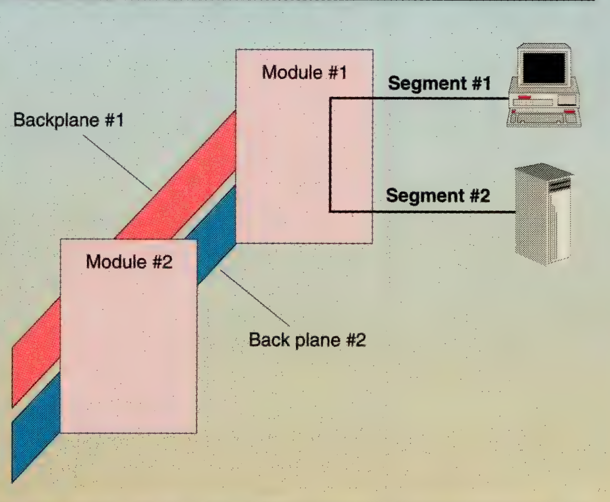

Figure 5b: Port Switching

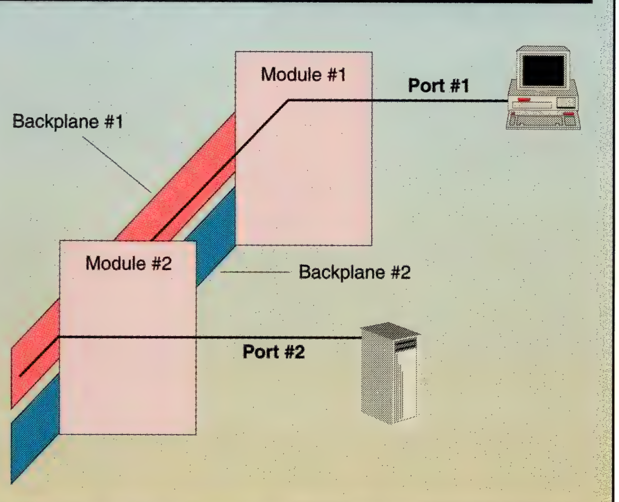

Figure 5c: Module Switching

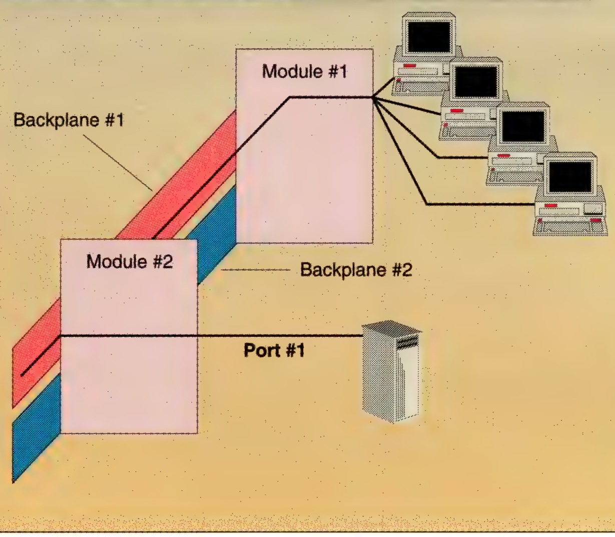

Figure 5d: Bank Switching

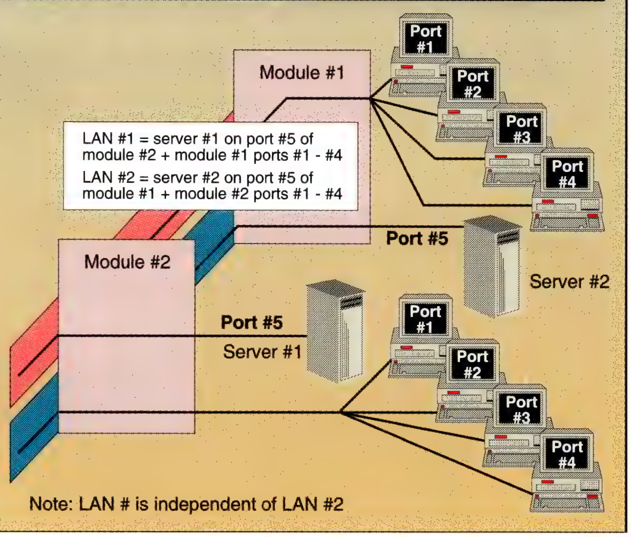

LAN #1 = server #1 on port #5 of module #2 + module #1 ports #1 - #4

LAN #2 = server #2 on port #5 of module #1 + module #2 ports #1 - #4

Note: LAN # is independent of LAN #2

to avoid the Management Module becoming lost as above.

Another important issue in deploying switched networks is the basis on which switching decisions are made. If a group of users is going to move then an administrator can manually reconfigure a network, but in situations such as unattended back-ups where the switching process must be automated, problems such as variable duration of a given server's back-up, a back-up application error, etc. must all be taken into account. How well this can be done, indeed whether it can be done at all, depends upon the level of integration between your network management environment and your operating system's environment.

At present those with environments based on more traditional hosts find this easier than those whose environment is based on PC servers. It is, for example, much easier to combine a TAR command, a command line SNMP utility and some error handling routines in a Unix shell script than, for example, do the equivalent under Windows NT or NetWare. This problem is historical in nature and will, hopefully, disappear as PC software vendors attempt to move into traditional systems vendors' markets.

In summary, switched networking is a very efficient, intelligent and cost effective technology. Its elements are very simple, but like a game of chess, their combinations can provide an infinite number of possibilities. Switched networks, be they Ethernet, Token Ring or FDDI are mature solutions — unlike the 100Mbps 'Ethernet' technologies under development (see 'Choosing Between High-Speed LAN Technologies,' in the February 1994 edition) — to the sort of problems which network managers are facing today. They are solutions which are standards compliant — unlike the 'full duplex' technologies which some vendors are testing in the marketplace.

Networks employing switched technology also provide great platforms for fault tolerant systems. Switched Ethernet also provides network administrators with the breathing space they need to await mature ATM products, an easier transition path to ATM when it arrives and experience in switched environments which will undoubtedly prove useful during a transition to ATM.

Why ATM? At present ATM is the only scalable high bandwidth technology on the horizon and it would appear that it will mature before, or at the same time as, the 'fast Ethernet' contenders or any as yet unknown high-speed network technology. As a result, those sites which outgrow switched Ethernet, Token Ring or FDDI will almost certainly have to go directly to ATM or face growing out of a 100Mbps Ethernet standard and *then* going to ATM. The longer this transition can be avoided the more mature ATM will be, the more options will be available and the more will be known about ATM's behaviour in real-life environments. After all, it's always best to learn from someone else's mistakes.

Graeme Le Roux is a Director of Moresdawn Pty Ltd (Bundanoon, NSW) and specialises in local area network consulting services.

The ACE ROUTER is the World's 1st fully Integrated routing solution for LAN Traffic over ISDN.

COMPLETELY INTEGRATED **Australian Made and Supported**

(No external Terminal Adaptor required)

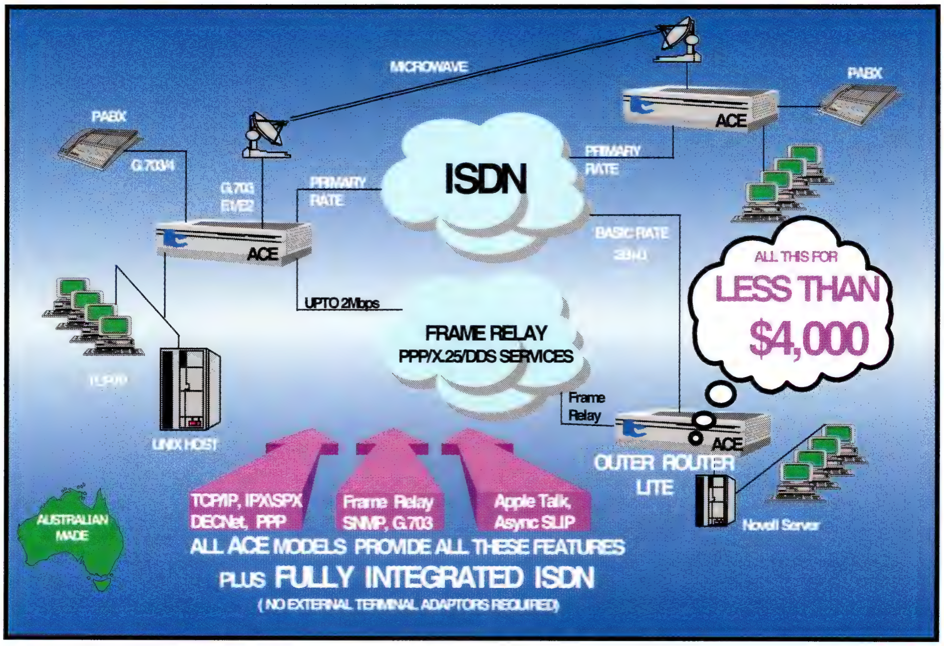

Interlink Communications

Head Office
PH: 02 975 2577
FAX: 02 452 5397

Melbourne Office
PH: 03 866 8913
FAX: 03 866 5387

Canberra Office
PH: 06 280 5368
FAX: 06 280 6383

Interactive Multimedia: The New Gold Rush

Data superhighways won't carry a whole lot of traffic until a new generation of multimedia applications appears. Stewart Fist examines the roadblocks slowing down development.

It's amazing who you meet on the Internet. I find that I'm in contact with a much wider range of opinion than I'd experience in real life; from right-wing American red-necks who believe in dynamiting the FCC and allowing AT&T, TCI and the Bell companies free rein to run the world, to watery-eyed world-weary wimps who believe all telephone systems should be public cooperatives run by environmentally-friendly social workers. You also make direct contact with researchers, technical experts and corporate executives at a level that is just not possible by non-electronic means. This expansion of horizons allows me to sit at home in Sydney and interview some of the best brains in this global business.

And they often provide colourful quotes. One of my correspondents characterised the recent telecommunications/broadcasting merger frenzy in the US as "a pack of hungry dogs with a sniffing curiosity for one another [which] has turned into an ongoing rut."

"The sweaty corporate copulating is a tad too frenzied," he maintained, and it will result in "multimedia mongrels and moguls roaming in packs as far as the eye can see." Good stuff! Great opening for a critical piece, I thought! But journalists need to be careful when quoting from Internet correspondence, because I later found out that this over-elegant turn of phase had been borrowed lock-stock-and-barrel from a speech by Wayne MacPhial, CEO of Southam Infolab. Which illustrates one of the potential problems we may have with data superhighways and the free exchange of information — what happens to copyright, fair-use, authentication of ownership, checking of sources, etc.?

MacPhial made his keynote speech to the US Information Industry Association's conference on Interactive Multimedia in Washington DC last September, and it's worth quoting some more (also from the Internet): "The bandwidth ballyhoo has everyone buzzing like high speed modems," complains MacPhial, who obviously spends his evenings dreaming up felicitous phraseology. "TCI and US West dream of a glass web ensnaring the couch-bound consumers of tomorrow; AT&T has us faxing our child to sleep from the beach; Philips still believes someone who doesn't work at Philips is interested in CD-I; and Barry Diller, the rhinestone cowboy of home shopping, imagines consumers shopping at Tiffany's over breakfast."

He is especially critical of the Bell Operating Companies, and their buy-up of cable networks and film libraries. But have you noticed that the buy-out business has recently gone quiet in the States; the realities of the bottom line have begun to sink in as companies learn to their cost that there's more to multimedia merger-mania than buying up your competition and suppliers. What they are now facing is the reality of a new convergent industry that needs to set some interactive multimedia standards.

Al Gore kicked off this 'superhighway' fanaticism when he compared an enlarged Internet with the American interstate highway system that his father helped to build. It was the 'vision-thing' of the 1960s, and the US highway network is often credited with

Continued on page 82

Standards and Technology

Have you been to a communications conference recently where there isn't at least a half-dozen incompatible desktop videoconferencing systems on show? All show a quarter-screen, indistinct image of a robotic nerd at a remote location who is keen to demonstrate that he can move the cursor on your screen and underline words in a memo — while periodically de- and re-materialising as he prattles on, with the all the audio quality of Donald Duck. It was the promise of this sort of 'Interactive Multimedia' that made ISDN such a raging success!

Which is not to say that interactive multimedia, either on-line or off-disk, won't be a success — it is just that the success or failure depends not on technical achievements, but on three essential factors, which are (to paraphrase the real-estate trade) 'standards, standards, and standards.' But probably because most of us are still trying to work out where to draw the line, the creation of multimedia standards is proceeding in a fairly haphazard fashion.

A leading would-be standards-maker is the Multimedia Communications Community of Interest (MCCOI) which recently held a forum in Sydney. The MCCOI is an international body established to develop internetworking standards for PC-based multimedia communications. Members include France Telecom, Deutsche Bundespost Telekom, BT, IBM, Intel, Telstra, BT and Northern Telecom. Members have pledged to work toward the establishment of open standards for multimedia internetworking and the development of universal functionality for desktop applications by adherence to an internationally recognised API. Sydney-based ISDN developer Jtec is Telstra's representative.

A number of companies have been attempting to establish standards for communications, command transmission and messaging over these new interactive networks. Microsoft is particularly active at a low level with its Microsoft At Work offering, but this has office applications only at present. The current leader appears to be General Magic, which was founded by two escapees from Apple: Bill Atkinson and Andy Hertzfeld (who jointly wrote most of the Mac operating system). General Magic has Apple, AT&T, Motorola and IBM backing, and Philips has recently taken a very strong position of support in Europe; it wants the European carriers to accept General Magic's approach as the de facto world standard.

At the present stage, General Magic's target is simply the personal messaging and electronic market services. Mid-term goals are to standardise communications at all levels between Personal Digital Assistants and desktop equipment, and its long-term intention is to standardise the signalling between all new multimedia devices and services.

General Magic's products and applications are based on the Telescript, an object-oriented programming language which has been specially designed for telecommunications, and the Magic Cap operating system, which contains a toolkit to implement comms functions and user interfaces.

Among existing standards, the MPEG-2 compression standard is really the main one which is internationally accepted. MPEG certainly isn't the first broadcast digital format by a long way. Digital video has been part of television studio operations for a couple of decades. What's more, since the mid-'80s there's been a Recommendation 601 for digital video in TV studios. R-601 (also known as 4:2:2) was devised by a CCIR committee headed by broadcast engineer, Jim Wilkinson.

The other multimedia standard which is widely recognised (but also often ignored) is H.261 which was devised by the ITU-T (ex-CCITT) for videophones and videoconferencing. It has a strong interactive component as part of a family of standards which now goes under the heading of H.320. Most manufacturers of videoconferencing equipment have their own proprietary compression techniques and control standards which they preferred to promote over H.261, but the H.320 family is generally provided as a fall-back.

Within the family you'll find: H.221, which defines the frame structure for telephone channels between 64Kbps and 1.92-Mbps; H.230 for frame synchronisation controls and indication signals; H.242, which provides the basic communications protocols; and a few others. G.722 is the standard for parallel 7 kHz audio coding, and there are image standards called CIF (Common Intermediate Format) and QCIF (Quarter CIF) now used as an intermediate stage between NTSC and PAL systems.

The H.320 family has proved itself to be an excellent videoconferencing standard, but one that has only limited application in the modern multimedia world. It faces the problem of being 'only' a conferencing standard, which means that the cost of investing in chip manufacture has not been spread over large production volumes, and so the cost of codecs remained artificially high.

MPEG and DVI

MPEG (Motion Picture Experts Group) is a family of standards which promises to bring costs down because of an enormous range of applications.

MPEG-1 was specifically designed for CD-ROM applications in competition with the proprietary favourite of the time called DVI (Digital Video Interactive), which had been developed by GE/RCA's Sarnoff Research Laboratories, and later sold to Intel. DVI was quite revolutionary for the time, but it has been modified, and now has possibly been superseded — although Intel doesn't see it this way. MPEG-1 was originally limited to data-rates of around 1.2Mbps, but this was later stretched to 3-5Mbps to accommodate higher speed CD-ROM discs, and some vendors have now released MPEG-1 products with data-rates up to 10Mbps. MPEG-1 was specified for a Source Input Format (SIF) image with a resolution of 352 x 240 pixels. MPEG-2 is backwardly compatible, and is designed for full 4:3 aspect ratio images of 720 x 480 pixels (and with true interlace scanning).

MPEG-2 was 'signed off' at the end of last year, and is being put though a form of 'conformance type-approval' stage. The MPEG group says that it 'cannot guarantee the standards are suitable for every application. [So they are] verifying the quality of the standard at the bit rates of 4, 6 and 9Mbps for a range of test sequences.' Preliminary results are already available, and the exercise should be complete by the end of 1994. What this means is that you can't be sure that one manufacturer's MPEG-2 decoder box will receive a channel encoded by another manufacturer's equipment.

Quite independently, the Digital Terrestrial Television Broadcasting project of the European Union has decided to standardise on MPEG-2 at three different data-rates. The 5Mbps rate will hopefully give PAL quality and permit reception by mobile units (in cars, etc.), while the 10Mbps version is for 'enhanced' TV, with data-rates of 25 to 30Mbps rates being reserved for high-definition television.

MPEG-3 was to have been a super high-quality MPEG standard for High Definition TV, but it has now been killed off. Instead, for HDTV, the operators will aggregate a number of MPEG-2++ channels. MPEG-2++ defines a technique of packetisation for data-rates in the range 15 to 60Mbps. MPEG-4, the latest in the bunch, is a very low-end standard for video-over-telecommunications networks running at rates of about 6Kbps. The idea is for videophones and conferencing systems to be run over ISDN B-channels, and still have bandwidth to spare for voice and cursor control. This standard is only in the early draft stages.

The key to DVI was a twin chip-set which ran at 12.5 mips. DVI images could be decompressed and displayed in real-time at resolutions up to 512 x 482 pixels, despite the very limited 150Kbps datastream from a CD-ROM. It was very clever in permitting 'layers' of image to be superimposed over the motion video — so text and graphics could be added to the primary picture for interactive teaching.

Currently DVI has two distinct algorithms: RTV (Real-Time Video) and PLV (Production Level Video). Both use a 4:2:2 digitised image as their primary source, but with a pixel-depth of 9-bits averaged over 4x4 pixel blocks. The DVI algorithms can handle image sizes up to 768 x 480 pixels, which is okay for PC-based applications.

Intel's new Indeo technology is a variation on RTV for desktop videoconferencing which is able to provide two small 160 x 120 pixel images side by side, but only at 12 frames a second. This is the current R2 version — one image is the incoming signal, and the other is the 'vanity' image which is out-going. The later R3 version promises 320 x 240 pixel screens (at 15 frames a second), but Intel doesn't specify what the actual image resolution will be — this could be a re-scaled figure.

DVI's main rival in the early days was a Philips development called CD-I (CD-Interactive). DVI differed from CD-I mainly in the fact that it was designed around standard Intel PC technology that couldn't successfully incorporate the OS-9 real-time operating system. CD-I was a standalone home-encyclopaedia system.

OS-9 has proved to be one of the most enduring and useful of all operating systems in this multimedia area. It functions in real-time, and it is behind CD-I and in many Motorola-based multimedia systems and has now been extended to the Intel platform. A later version written in C is known as OS-9000. What appears to make OS-9 so useful is that it has a shell and a scripting language similar to the Bourne shell found in Unix. Files are in a tree structure with separate user-areas, and everything, including devices, is treated as a file. Don't ask me why, but apparently this is important.

DVI replaced OS-9 with its own 'real-time executive' to handle audio/video/disc access concurrency. Intel has since developed a special Intel 750 chipset for DVI compression and decompression, but it can probably only succeed against the open MPEG standards if the company manages to make the chipsets cheap enough.

Intel says that its Indeo videoconferencing unit will sell for less than $US2,000 and run over ISDN, analogue phone lines or local area networks. This appears to be a tall order. Indeo is also the basis for Intel's Smart Video Recorder which is bundled with Microsoft's Video for Windows at $US700 in the States. This system is used to capture full-motion video.

While the video compression scene is in turmoil, vendors are often taking a safe course by defining 'frameworks' or 'pigeon-holes' rather than proprietary decompression systems. These are frameworks for plug-in decompression procedures, and they include the intelligence to select and activate the most appropriate algorithm.

MPEG-2 happens to supply such a minimalist 'framework' as does Microsoft's Video for Windows, Apple's QuickTime, and IBM's Multimedia Presentation Manager/2 product.

To the Desktop

The race to bring these technologies to the desktop has already been won and every backyard solder-shop capable of mounting computer chips onto a PC board has something similar on the drawing board.

Silicon Graphics, Sun, DEC and Sony all support videoconferencing as an off-the-shelf feature in their new workstation releases. As you'd expect, they've charged ahead using compression standards — but in a way that ensures that their workstations will lack interoperability or mutual connectivity across vendor product lines.

To get anything to work, you will need to buy all of your workstations from the same vendor, and link them internally across LANs or private networks with some brand new and expensive equipment. Ethernet, with its maximum transfer rate of 10-Mbps, gets bogged down even when transferring voice memo files. This is why the workstation manufacturers assume that 'ATM' is indeed an acronym for a cash-dispensing protocol — we will need the promised ATM data transfer rates of 34 and 155Mbps before we can begin to take desktop conferencing seriously.

Recent trials of Video Internet as a TCP/IP-based wide area multimedia application was a curiosity, but it doesn't seem to have a future. Internet protocols are unable to guarantee delivery, and TCP/IP needs a mechanism for multimedia (voice and video) synchronisation.

Delays across data networks are a constant problem. Sony has decided to 'time-stamp' its video data to resolve jitter (delay dispersion) problems. It's a good idea, but currently it has no wide acceptance. The transmitting workstation adds a high-precision time-stamp to each data packet, and the receiver can then buffer the data, and release it with the correct timing.

Sony's time-stamping approach also incorporates network mechanisms to turn up, or down, the rate of transmission so that the video connection comes under self-control when traffic loads peak. While this variation messes up the image quality it does make stable images possible and reduces the demand on network resources, because it limits the number of network collisions.

Some researchers are working on new multimedia networking protocols called ST-II and TP++, but these share no significant relationship with current network protocols and don't offer a migration path to new developments. A new approach appears to be essential if video is to be transmitted over Ethernet (even Fast Ethernet) or Token Ring systems where irregular delivery of video packets is the norm.

However, any process of this kind generally means that you'll be watching images which have been delayed for half-a-second or more in buffers along the way.

The Japanese appear to have fallen well behind in their use of compression techniques. Currently Sony only provides JPEG as the compression system, although it has a H.261 standard compression card under development. Sun, by contrast, sells the SunVideo card which can handle MPEG-1, JPEG, H.261 or its own proprietary compression standard called Cell.

Silicon Graphics' new desktop workstation (called Indy) has all of these compression systems as a standard feature, and it adds Apple's QuickTime to the list also. But Silicon Graphics promotes MVC1 because this is its own proprietary standard.

DEC offers an Alpha AXP card with MPEG-1, JPEG, QuickTime, Video for Windows and its own SMP proprietary standard — but no H.261. Both DEC and Silicon Graphics have proprietary video mail and the shared whiteboard functions, along with their videoconferencing.

The new workstation systems differ in other ways also. Some are direct-connected to a small video camera, while others can replay video information from disk or CD-ROM. DEC and Silicon Graphics have licensed Apple's QuickTime video formats which are widely used on the Macintosh, and DEC is still working on providing Video for Windows in its workstations.

At a minimum, what is needed before all these workstations can interconnect is:

- An API specifying the compression, decompression/play function;
- Defined data formats (including the compression scheme); and
- A services protocol for the network (between client and server).

Many vendors are members of the Interactive Multimedia Association, and they say they will abide by the IMA's standards when they are established.

All said and done, it's early days yet!

Stewart Fist

From page 79

being the major public project after the war responsible for much of America's later affluence. Gore believes that a glass-fibre Infobahn is the 'vision-thing' for the Clinton-Gore information society of the 1990s. And he's probably right, although he may be out by a decade or so.

Superhighway may be a relatively new buzzword, but 'Broadband,' 'Interactive' and 'Multimedia' have been around for some considerable time. Broadband has always meant television delivery, and both Interactive and Multimedia were part of the lexicon of 1980s educationalists who fondly imagined that life would be easier when computers could take over the hard-slog of repetitive teaching.

'Interactive Multimedia' meant to them, 'self-paced learning' — using VCRs and videodisc players with some sort of computer control, as a user-controlled teaching machine. Later it came to mean videogames. (Which, come to think of it, is much the same thing — except that the kids probably learn more from the games!)

But in the last six months:

- *Multimedia* has ceased to be a local concept linked to CDs and VCRs. It has now absorbed the concept of 'convergence' and has metamorphosed by incorporating global networking and broadcasting.
- *Interactivity*, which we previously applied to local control systems using computer selection, has migrated to mean remote access and sophisticated control over network servers and broadcast (one-way) systems. Broadcasting has come into this new 'interactive' domain through convergence with telephony (two-way) networks.
- *Superhighway* has become synonymous with universal to-the-home two-way video. It originally meant interactivity between university workstations and remote shared supercomputers over optical fibre links. Later it came to mean the activation of dark-fibres available between the various exchanges, reaching out to many primary server nodes. Now it seems to refer to the home and small-business access network. With this degree of semantic migration, you can safely apply it to anything faster than a 300 baud modem link.
- *Convergence*, of course, is the key concept here. I would suggest that, in the technical sense, it just means local computing, telecommunications and broadcasting systems interacting — and that's not very new. But for the telephone companies, technical convergence represents a chance to break out of the gold-plated strictures of just being a carrier of voice and data and make real money from entertainment broadcasting and interactive services like home shopping.

From the other side of the fence, for broadcasters it's a chance to get their fingers into the lucrative market of telephone carriage. The gold always seems more plentiful where the other fellow has staked his claim.

Over and above all this is an economic theory similar to that of economies of scale. But here it is the 'profitability of dominance,' and there's little doubt that whoever controls the broadband version of the interactive local loop will be able to dictate terms to all the other players.

You can see this infectious profitability concept spreading like mad-cow disease through the brains of the cud-chewing executives in America. Bell Atlantic apparently believes from its surveys that the per-household income from interactive multimedia will be $US1,745 ($A2,500) in the first year of ADSL, and they say that by the end of the century each home will spend $US9,827 ($A14,260 in real-dollar terms) of their disposable income on on-line games, home-shopping and entertainment. I suspect that this figure was actually the zip code of the researcher's love interest at the time — or maybe they slipped a decimal.

Bell Atlantic isn't alone in its hallucinogenic optimism; the big-figure disease has taken on epidemic proportions. Last year a major North American cable company was sold for a figure which represents $5,000 per household connected because interactivity was just over the horizon. Yet the average income per household connected by cable (and remember this includes cash passed onto the program producers/distributors) only averaged $US300 p.a. Before the lithium tablets took effect, Time-Warner executives were planning to spend $US750 million on 4,000 Orlando, Florida homes for interactive multimedia trial purposes ($US187,500 per home). When sanity finally got a grip on them, they cancelled the project.

For the rather innocent television program makers and games software writers, convergence presents a lure of 500 channels — with little realisation that this doesn't necessarily mean more people are watching and playing. The key principle to remember here is that of 'dilution.' When you have 500 channels, you also have 500 programmers — all of whom will be running their best program in the same peak-viewing timeslot. More actually means less; you will miss all the good stuff you wanted to watch unless you've got 500 programmable VCRs.

For agents and program distributors, 500 channels means many more distribution deals. More companies will be bidding for the best (or indeed 'any') product, prices for anything watchable will skyrocket, and naturally the distributors will expect to get a bigger slice of the cake because of their increased workload.

For those educationalists who worship the God of Technology, it means the Utopian vision of better amortisation of resources, distance learning for all, self-paced training, and all the other patina of promises that they think will reduce their workload and solve the problem of school boredom. As has happened so often in the past, they will probably be sorely disappointed. And for business, interactivity means new ways to generate responses and to contact customers/clients . . . and so on. In fact, there's a parcel in this basket of interactive multimedia goodies for just about everyone.

But convergence brings problems not just advantages. And the main problem it introduces is one of defining standards (see 'Standards and Technology' on page 80).

A Standard Dilemma

There's an incredible naivety in the standards area, generated by the euphoria of the personal computer boom. Everyone wants to do a Microsoft: racing out to buy a Quick & Dirty OS and flog it next day to IBM as PC-DOS, thereby instantly creating a de facto standard.

But standards aren't so easily made in broadcasting or telecommunications, and many of the players don't understand that technical convergence also brings with it a convergence in industry cultures. Computing, broadcasting and telecommunications industries look at standard definition in quite different ways.

In the past, our computer systems essentially presented only *local* standards problems. We could afford to use vendor-specific solutions if the applications, controls or data-formats were confined to one machine or to one local network. Eventually some wider standards evolved by long-term attrition, and through market dominance. CD-ROM is a form of Interactive Multimedia from this culture.

Broadcasting, of course, requires *national* standards. You don't really want Perth using SECAM television, Hobart on NTSC, while the rest of the country watches PAL. Nor do you want your satellite raining down digital wide-screen, while your terrestrial viewing is by normal PAL — although with Pay and satellite systems you do at least have a new box in the delivery chain, which can be called upon to perform standards-conversion.

Telecommunications has an even wider standards emphasis than broadcasting: it must be *international* in scope. The phone on my desk will make contact with a payphone in rural Mongolia; I can Telex messages to a hotel anywhere in the world; and my access to mainframe computers via globally interconnecting X.25 packet networks is almost universal.

In fact the great success story of telecommunications for the last hundred years, has been its ability to retain international standardisation during times of ideological aggression, hot and cold wars, political antagonism, and nationalistic turmoil. But one of the questions we certainly must address is whether this standardisation can or

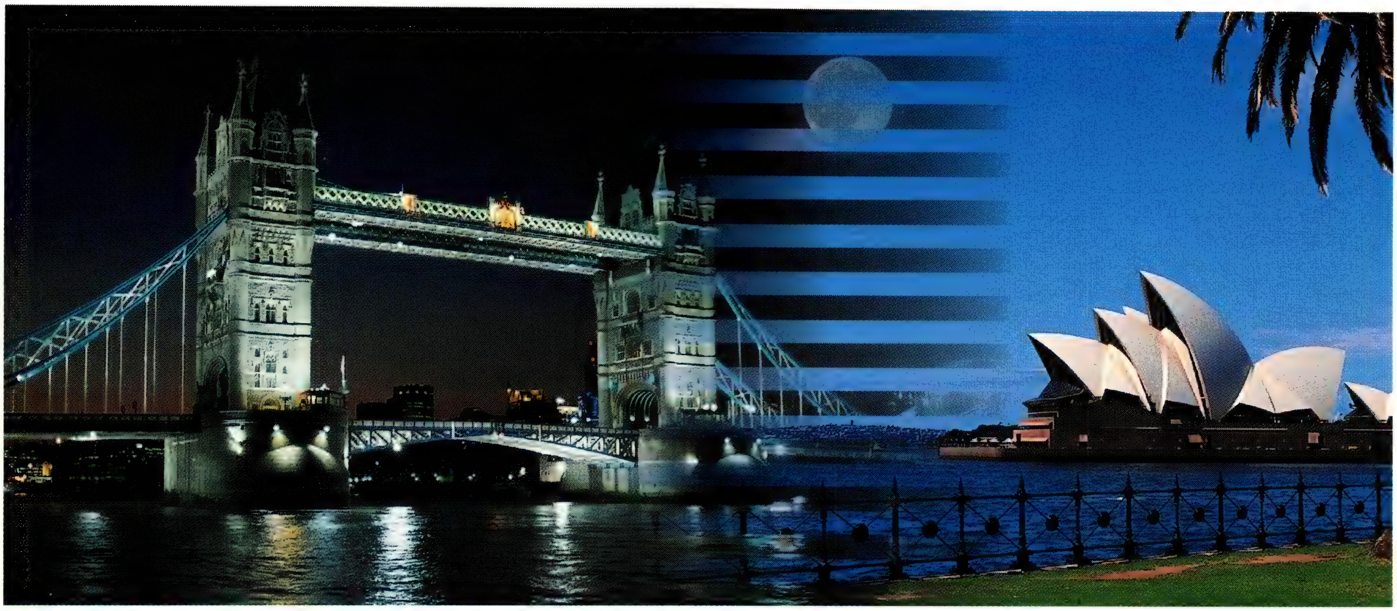

BRANCH *office* CONNECTIVITY *is* FINALLY *within* REACH

Until today the prospect of integrating a branch office into your corporate network has always seemed hopelessly far off. There are now unparalleled connectivity solutions that bring your branches and headquarters much closer together.

Eicon Technology, a world leader in PC-based connectivity, offers a family of internetworking and host access products for linking remote offices, along with some of the best 3270, 5250 and APPC emulations for use on your favorite PC platform. Eicon's new InterConnect Server extends your Cisco, Wellfleet or IBM backbone network across the enterprise. It fuses traditional router and gateway technologies by integrating LAN interconnection for Novell NetWare and TCP/IP networks, access to any host, as well as conversion and concentration of non-LAN communications—all supported by comprehensive SNMP network management.

If you are interested in connecting your branch offices—whether they are in London, Sydney, or anywhere else in the world—call us or send the adjacent coupon for your Branch Office Connectivity booklet.

Please send this coupon to:
Eicon Technology Corporation
Level 14, 33 Berry Street
North Sydney NSW 2060
Australia
Tel: (02) 959 1960
Fax: (02) 959 1980

ACM (AU)

NAME:

POSITION:

COMPANY:

ADDRESS:

POST CODE: COUNTRY:

TEL: FAX:

New Australian Office

EICON
TECHNOLOGY

will continue to develop in a true competitive environment. And the real test may not be between telcos, but between the future multimedia program providers.

"I worry that we may soon see three or four omnipotent media giants building information highways that, for all intents and purposes will remain one-way streets," says MacPhial. And all the evidence of recent months shows that he has cause to worry.

One-way streets do, however, largely solve the standards problems. You don't need to decide which side of the road on which to drive in one-way streets, and, if there are only 'three or four omnipotent' companies, they can collude in deciding what they will provide and at what price. Dictatorships and oligopolies have never had problems with establishing standards.

But as MacPhial says, "when all the smoke clears from all the mirrors we'll be left with a few large, inbred companies providing a lame, limited version of interactive multimedia consisting of home shopping and networked games.

"Joan Rivers has just been signed to do a program called *Can We Shop?*," he points out, "and [Arnold] Schwarzenegger shot extra scenes for the *Last Action Hero* so he can produce an interactive version of a linear bomb. These are not good signs. The future of multimedia isn't about one-to-many. It's about many-to-many. It's about building communities."

I'm not sure that I'd dignify the flaming that goes on over the Internet as 'building communities,' but I certainly agree that we've got to view this interactive multimedia future as more than just an avenue for big business to feed more junk to more consumers over wider global networks.

And it always seems apparent to me that interactive multimedia, convergence and superhighways are more about telecommunications than about broadcasting or computers. So I think it is obvious that we need the standards-definition skills of an international forum like the ITU to take over the definition of multimedia and interactivity standard — even if it slows progress.

If we examine the current state of interactive multimedia, we really have only one substantial standard which is internationally accepted. This is the MPEG-2 compression standard, which has passed through the last definition stage, but is yet to appear in silicon — although there are a few 'near-enough' implementations.

As Ripley said: It is strange but true. The only global standard that we've managed to devise for the broadband superhighway is one that would probably be unnecessary if we had broadband superhighways. MPEG, like all compression standards, is designed specifically to cram a wide truck-loads of data down narrow data-lanes.

You only really need compression to compensate for inadequacies in the network. And these are not necessarily technical inadequacies — the main defect in Australia is in the *cost* of providing bandwidth. Currently these prices are related to what the customer will bear, not to the real cost of bandwidth provision.

Telecom and Optus now have a combined interstate and interexchange bandwidth which is almost unlimited in practice, if they bother to fire up their dark fibres. As an example; Australia has about 300Gbps of bandwidth for just the existing Sydney/Melbourne optical fibre links (not counting the microwave and satellites), which is about 64Kbps for every resident in both cities, 24 hours a day. Do you want a data-superhighway between the States? You've got one! Now just work out how to pay for it at these ridiculous long-distance prices.

The Gold Rush

Of course what will really drive multimedia technology is a killer application which gives people a reason to buy — like Lotus 1-2-3 did for the PC. On the desktop, AT&T, MCI and BT released videoconferencing units last year in the form of videophones, but these were analogue devices that could not talk to each other. And to suggest that they have been a dismal failure would be a serious understatement. The carriers can't seem to give them away.

Videophones have long been central to our the vision of the broadband future, but you may have noticed that, as the superhighway concept has grown, the emphasis on videophones has diminished — almost disappeared. Long-term fibre-to-the-home trials consistently show that customers rate home and small-business videophones as tiresome trivia. Multimedia telephony appears only to work in narrow areas of business and academic research — but the manufacturers won't believe their own surveys. They still have multimedia mass-markets in view, and they are blinded by the current epidemic of optimism.

Logic doesn't seem to enter much into these billion dollar development decisions. Forrester Research, for instance, has released a prediction that it would be the year 2000 before many interactive multimedia products returned their development costs, but that doesn't seem to worry investors.

And the entertainment side of multimedia investment is not necessarily more lucrative either. People seem to forget that there are now 11,000 cable companies (six large ones) in the United States, and few of them make more than average industrial profit margins — even though they generally hold a local or regional monopoly. Like Rumplestiltskin they've been busy weaving straw for years, but so far there's precious little gold in evidence.

The breakthrough with the most immediate profit potential, according to those who seek to bridge the broadcast/interactive gap, is in set-top boxes which can receive addressed messages, and send out coded responses. Cable or over-the-air television can then become truly 'interactive' — but not in the sense of videophones.

This has all been done before, of course. The People Meter in a few thousand Australian homes already has this sort of two-way function, and many years ago Qube in Columbus, Ohio, had a full two-way 'response unit' that enabled you to register opinion, vote, select special services, and ask for individual programs.

Later a system called 'Consensor' was trialled in Alaska. It improved on Qube by permitting you to register not only priority, but intensity of feeling through the manipulation of two dials. Consensor became an experimental electronic plebiscite system for local and state government; residents would register their opinions after listening to a background documentary/discussion on such issues as port and road improvements, jails, mass transit, and exploration of natural gas. The results on these issues could be flashed on the home TV screen within one minute. Both Qube and Consensor died quietly many years ago: Michael Rimmer was right.

There's always a lot of talk about interactive education and cultural uses of the data-highways, but, as Paul Saffo at the Institute for the Future (California) says: "The dirty little secret in the information revolution is that no-one is really interested in information." Game-playing, virtual-reality, shopping, theatre booking systems and electronic punting may be less worthy uses of interactivity, but many investors think they are likely to be more profitable.

Interactive Television Australia certainly sees it this way. Managing director Jim McKay has bought forward trials of his Australian-developed set-top box to August of this year in a race to avoid being swamped by dozens of American competitors already racing to market. Interactive Television will place 1,500 boxes in an Adelaide trial to prove the concept before a stock-market float. The boxes allow couch potatoes to participate in game shows, have a bet, or choose a tacky Xmas present from a home-shopping program without stirring from their semi-stupor.

Interactive's box has a small printer, and it uses the telephone lines to zap information back to base. It is designed to work with analogue or digital signals, and it can be applied just as well to Australis-Lenfest's satellite/MDS system or Telecom's coaxial — or the current UHF/VHF networks. McKay says the company will give 200,000 boxes away at the national launch later this year, so he certainly expects to clean up in terms of his on-line activity.

Stewart Fist is a freelance journalist based in Lindfield (NSW).

Finally, a hub you can wrap your whole network around.

Introducing the Lattis System 5000.™ The first network center hub.

Think of it as the hub of hubs.

SynOptics' new Lattis System 5000 is the first intelligent hub to function at the center of your network. Instead of simply connecting users together within a LAN segment, the System 5000 allows you to build a structured network, linking dozens of LAN segments to the network center. The result is the best of both worlds: the security of a central location to consolidate and manage your key resources, and the flexibility of distributed LAN segments.

The LattisSphere System.
Where hardware and software work together for a higher level of control.

Each System 5000 supports up to 52 Ethernet segments, 26 Token Rings or 28 FDDI paths, connecting tens of thousands of users. And it has a built-in migration path to future technologies like ATM and high-speed LAN switching.

Fundamental to the structured network is SynOptics' new Optivity™ software, which takes an entirely new approach to network management. With its intuitive user interface and a simple drag-and-drop operation, you can assign network devices to logical workgroups—regardless of their physical location.

To learn more about the Lattis System 5000 and SynOptics' new LattisSphere™ System, call Australia 1800 81 7070 or New Zealand 0800 44 7070. We'll be happy to send you the complete wrap-up.

SynOptics

The Network Fabric of Computing

© 1994 SynOptics Communications, Inc. Lattis System 5000, LattisSphere and Optivity are trademarks of SynOptics Communications, Inc.

CALL AUSTRALIA 1800 81 7070 or NEW ZEALAND 0800 44 7070.

Telecom/Datacom Testing with PFA-30, The Complete Solution

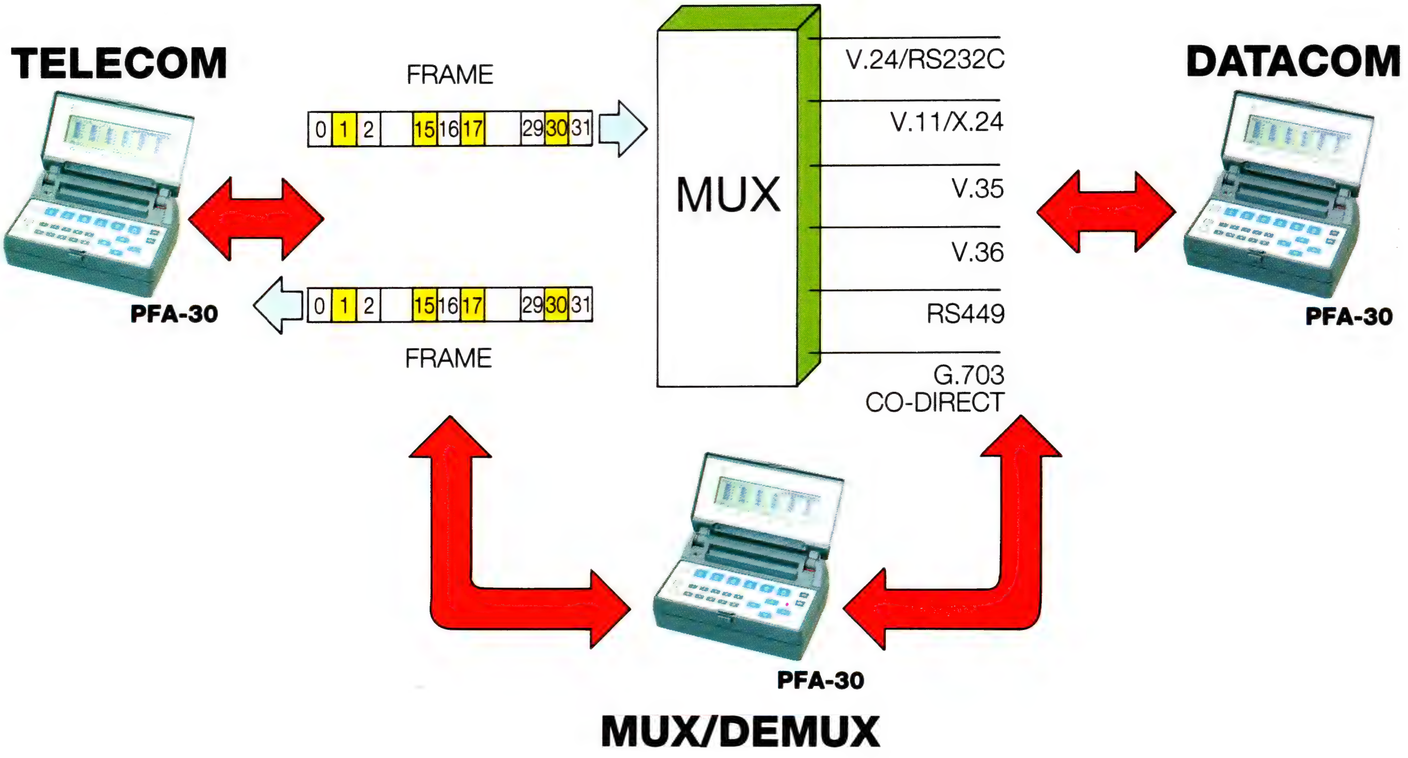

TELECOM/DATACOM TESTING

The increasing variety of interfaces and equipment used in modern voice and data networks creates a need for versatile and comprehensive test equipment. The emphasis is on portability, for field use, together with the ability to test all the interfaces likely to be encountered.

PFA-30 combines telecom and datacom testing capability in a small hand-held package.

THE COMPLETE SOLUTION

The combination of Multi-Interface working, Multi-bitrate clock generation, DCE/DTE emulation and MUX/DEMUX testing in a single box enables the PFA-30 to provide the complete solution for data service providers, maintenance technicians and end-users.

Wandel & Goltermann Pty Ltd
42 Clarendon Street
South Melbourne Victoria 3205
Tel: (03) 690 6700
Fax: (03) 690 6750

Wandel & Goltermann
Electronic Measurement Technology

Bringing ADSL Home — The Race is On

ADSL may hold great promise but can it be fulfilled? Robin Whittle examines the practicality of rolling out this ambitious technology fast enough to meet the challenge posed by fibre optic and coaxial cable.

In the last two years, Asymmetrical Digital Subscriber Line (ADSL) technology has come from nowhere to be a promising contender in the race to provide individual high-speed interactive connections to each home, office, factory and school. ADSL transforms existing twisted pair subscriber loops into digital links capable of supporting up to 6Mbps downstream from the telco exchange to the home and 640Kbps duplex between the home and the exchange. It performs this miracle — which may lead to the early introduction of Video on Demand (VOD) and other interactive subscriber services — whilst allowing the POTS phone service to operate normally (see Part #1 of this article, 'ADSL — Bridging the Superhighway Gap?,' in last month's edition).

But it's still early days for ADSL and in its current pre-standardised form the technology is very expensive and relatively immature. So for this revolutionary technology to fulfill its promise and move out of the lab and into service many questions of cost and performance will have to be resolved. While sceptics abound, recent developments point to approximate answers to these questions and indicate that ADSL will probably play a significant role in our telecommunications future.

How Far Will it Reach?

The question of DMT (Discrete Multi-Tone) ADSL's effectiveness could be reduced to 'What proportion of the population can receive how many megabits of downstream data, at various duplex data rates, with acceptably low error rates?' This depends on the distribution of twisted pair cable lengths, and how DMT performs on those cables.

The cable length question can be answered approximately by considering the current design guidelines for installing twisted pair POTS services and by referring to the most recent survey of Telecom Australia's Customer Access Network. The survey was conducted in 1979 and 1982 and covered 719 urban subscriber loops. 77% of loops were 3km or less, 94% 4km or less and 98% 5km or less. The great majority of these loops used 0.4mm copper cable insulated with either polyethylene or paper. 44% had one bridge tap and 7% had two or more. Bridge tap lengths ranged from 10 metres to 4km, with an average of 162 metres. Only 5% of loops ran straight from the exchange to the home. The rest went through one or more pillars or cabinets. The average number of cable type changes was three, but 2% had eight or more.

Telecom guidelines for POTS services state that the attenuation at 820 Hz must be no more than 6.5dB. This sets certain distance limits for subscriber loops used by the great majority of the Australian population.

Table 1 on page 89 shows reach estimates based on computer models and actual performance of leading ADSL developer Amati Communications' Overture and Prelude transceiver products. Amati chief, John Cioffi, says that the 3.7km figure is based on a computer model for the still to be released Overture and takes into account a number of loop topologies, with no interference from

Who is an *Australian Communications* Reader?

- There's a better than even chance that they work for a company with over 500 employees
- There's a 70% chance that they authorise or specify the purchase of data networking products
- It's even money that they run a Novell network or will purchase one in the next 12 months
- There's a 70% chance that they are running Windows or will be in the next 12 months
- And there's an almost 80% chance that they prefer *Australian Communications* for depth of information and value to them in their jobs over any other Australian IT publication

And there are 26,500 of them who read *Australian Communications* every month. Maybe you should be talking to them!

***All statistics are based on the response to the August 1993 *Australian Communications* readership survey conducted by Price Waterhouse, and *Australian Communications'* Audited Paid Circulation.**

For a copy of the complete survey results please contact Craig Burkill or Dawn Switzer on (02) 264 2200.

Price Waterhouse

Table 1: ADSL Reach Limits

	COPPER CABLE DIAMETER		
	0.4MM	**0.64MM**	**0.90MM**
Distance limit for POTS	4.2km	7km	11km
6Mbps Amati Overture estimate with 6.144Mbps simplex + 640Kbps duplex	3.7km	6.2km	9.7km
6Mbps Telecom Research Labs estimate with 6.144Mbps simplex + 448Kbps duplex	3.5km	6km	7km
2Mbps Telecom Research Labs estimate with 2.048Mbps simplex + 176Kbps duplex	>4.2km	>7km	>11km
6Mbps Amati Prelude performance with 16Kbps duplex	2.7km	N/A	N/A
2Mbps Amati Prelude performance with 16Kbps duplex	4.0km	N/A	N/A

Source: Telecom Research Laboratories, Amati

other services such as HDSL or T1 — which is a major problem in the US. The Telecom Research Laboratories estimates were calculated for 0.4mm cable and extrapolated to the other gauges. The 2Mbps figure is not available but was well above the 4.2km POTs limit.

The TRL estimates considered downstream data at 2.2 and 6.3Mbps, starting with sub-carrier 15. This is not realistic for Frequency Division Multiplexing (FDM), which should start with sub-carriers well above 17. For echo-cancellation, it leaves sub-carriers 5-14 to be used for perhaps another 440Kbps downstream and with a similar or greater rate upstream — depending on the limitations imposed by NEXT (see 'Inside DMT' on page 90).

For both the TRL and Amati models the noise margin for a BER of 10^{-7} was a conservative 6dB. The TRL figures were derived from worst-case crosstalk models with two E1 or four 2B1Q HDSL disturbers in the binder group, with the remaining pairs carrying DMT ADSL.

TRL's Phil Potter says that there should be little interference from Basic Rate ISDN signals — either the 4B3T of the Alcatel NT-1s which are installed now, or the 2B1Q signals likely to be used in the future. However ISDN links could be replaced with ADSL if they caused excessive interference. E1 lines interfere more strongly with ADSL, but are less common. They could be replaced with HDSL (High-bit-rate Digital Subscriber Line) — which causes less interference. However, Telecom has decided that the fairly modest cost reductions offered by HDSL do not justify its introduction at present.

None of these models or tests include the effects of impulse noise from sources such as refrigerators and ringing signals which are interrupted when the telephone is answered. Little is known about this and Telecom is surveying the network at present. Neither do these models include the wiring to the home. Many homes are connected via an overhead figure 8 'drop wire' and their attenuation and susceptibility to noise has not been studied.

If the computer models can be believed, then DMT ADSL could provide 6Mbps on at least 90% of subscriber loops, and 2Mbps to almost every home within 11km or so from their exchange — as the cable lies. However, the models could easily be unrealistic. Crosstalk is known theoretically but the real test is a cable full of DMT signals. A practical POTS splitter, hybrid, filter and ADC (Analogue-to- Digital Converter) all introduce noise and distortion — which probably accounts for the differences between the computer estimates and the actual performance of the Amati Prelude.

Another reason to suspect that these models are optimistic is that a recent Amati contribution to the ANSI ADSL standards committee gave a 6Mbps rate for 4.2km of 0.4mm wire and 8Mbps for 3.7km with no HDSL crosstalk and an unspecified duplex channel. But this was for a non-standard system with sub-carriers beyond 1.1 MHz and with double the usual signal power.

This double DMT system, which has downstream sub-carriers as high as 2.2 MHz and four times the number of upstream sub-carriers, is intended to give 15-Mbps downstream and 1.5Mbps upstream in the presence of T1 noise, but only for about 1.5km.

Telecom may not want to offer a service it cannot provide to almost everyone and several things must be done before it can be confident about providing 6Mbps ADSL to, say, 99% of the urban population in Australia. Extensive field trials of high performance DMT transceivers must take place. Impulse noise and drop wires must be studied. A new survey of existing subscriber loops must be conducted to determine what rates of ADSL they can support and what it would cost to bring 6Mbps to homes with difficult loops.

Cost and Availability

The cost side of the ADSL cost-effectiveness question is far from settled and is dependent on efforts to integrate the transceiver into a chipset.

All DMT transceivers currently used in field trials are Amati Communications Preludes which operate without trellis encoding and use FDM instead of echo cancellation. They are based on a number of Motorola DSP56000 24-bit DSP CPUs placed on boards inside a box a little smaller than a mini-tower personal computer. They cost several tens of thousands of dollars a pair, and provide up to 6.176Mbps downstream and 16Kbps duplex.

Motorola is now working together with Amati to integrate the major analogue and digital functions of a DMT transceiver onto a single chip and sell them for less than $US100 each by the first quarter of 1996.

Continued on page 93

DMT Standardisation Efforts

The ANSI T1E1.4 committee has been working toward an Issue #1 DMT standard for release this month. This will specify a basic interface to CPE (Customer Premises Equipment) which will suffice for the large scale trials now planned over the next year or so.

Before DMT ADSL can be deployed widely there needs to be a standardised, more flexible interface to CPE, so that users can plug in various items of equipment such as set-tops, phones and computers. This interface is expected to be a multi-drop bus where one or more items of CPE can be daisy chained with short simple cables — like the S-Bus which connects between the ISDN NT-1 box and phones, faxes or computers.

The design of this interface is likely to be contentious. One feature it will have which is not supported in the Issue #1 standard are flags to warn the CPE that a particular block of data is known to be erroneous. Perhaps this interface should be based on ATM.

A mature CPE bus which can be used for mass produced DMT transceivers is the major addition planned for Issue #2 in a year's time. By then, considerable experience will have been gained from design, manufacturing and field trials.

Robin Whittle

Inside DMT

In an ADSL link to the home, digital data is sent in both directions using frequencies above 20 kHz so as not to interfere with the POTS service which shares the line. The most difficult task is to send and receive up to 6Mbps of downstream data from the exchange to the home. Upstream data is conveyed in a similar way using fewer sub-carriers. Each transceiver (at the exchange and at the home) uses powerful DSP hardware to perform the transmission and reception operations.

This is a complex process which was originally achieved by AT&T with a CAP (Carrierless Amplitude/Modulation) single carrier technique. CAP has now been overtaken by DMT (Discrete Multi-Tone) which is currently being standardised by ANSI's T1E1.4 committee.

DMT adapts its coding to make the most of the signal-to-noise ratio available at differing frequencies. The various logical channels of user data are aggregated with the management data (the transceivers are constantly comparing notes) and broken into 4,000 frames per second. Each frame could contain up to 1,728 bits. Data which is urgent and is more tolerant of errors is assigned to a fast data buffer — this would include voice telephone audio data which should not be delayed by more than 2 msec. Non-urgent data — including any which must be protected from errors — is sent to the interleave buffer.

Each buffer has a CRC (Cyclic Redundancy Check) byte appended to it and the contents of each buffer are fed through a scrambler which inverts about half the bits according to a predictable pseudo-random sequence. A synchronised descrambler at the receiver flips them back to restore the original data. This does not increase the number of bits, but breaks up any continuous run of ones or zeros — which could degrade the performance of the system by causing long periods of high sine or cosine levels, instead of the randomly distributed levels which result from scrambled data. Many modems use scramblers for the same purpose — that is why they make a noisy sound even when transmitting no data.

Both buffers are then processed for FEC (Forward Error Correction) with the Reed-Solomon algorithm which takes a set of bytes and appends to them a number of FEC bytes such that if any one byte in the set is corrupted, it can be restored by the receiver. Multiple byte errors can be detected but not corrected.

The resulting bytes in the fast buffer are sent in the current frame and those in the interleave buffer are sent one at a time in successive frames. The idea is that a spike of impulse noise may corrupt a whole frame, but this will only change one byte of the set from the interleaved buffer. After the bytes have been reassembled in the receiver, the Reed-Solomon correction algorithm will be able to find and correct the corrupted byte. Reed-Solomon codes are extremely complicated, very well accepted and highly efficient.

The Reed-Solomon code adds between 4 and 8% to the number of bits to be transmitted and the interleaving increases the latency of the data to a maximum of 17 milliseconds. The combination of the two greatly reduces the error rate in the presence of impulse noise. DMT ADSL uses one level of Reed-Solomon and interleaving while the Compact Disc and DAT use two.

Some sub-carriers may carry as little as two bits with 4 QAM and others may carry up to 12, and so have as many as 4,096 constellation points (see 'ADSL — Bridging the Superhighway Gap?' in last month's edition for a description of QAM (Quadrature Amplitude Modulation) data transmission techniques). Amati's John Cioffi says that up to 15 bits can be encoded on each carrier while Marcos Tzannes of Aware says that they are unlikely to encode more than 11. During training — when they are turned on or after a major disruption — the transceivers test the loop and determine how many bits of data can be can be reliably transmitted on each QAM sub-carrier. During operation, they monitor loop conditions and may change the bit assignments to optimise performance.

Without trellis coding, the data bits directly select a point in the QAM constellation and this determines the level of sine and cosine wave for this carrier. If trellis coding is used, then the algorithm is much more complicated.

Trellis coding in the transmitter, when combined with a Viterbi decoding algorithm in the receiver, can be used to improve the system's performance in the presence of continuous noise. It reduces the bit error rate and/or increases the number of bits carried — without increasing the bandwidth or average energy of the sub-carriers. Like Reed-Solomon, it is complex, widely used and impossible to explain briefly. Trellis encoding, Viterbi decoding, interleaving and Reed/Solomon encoding and decoding are onerous tasks with general purpose DSP chips but can be performed efficiently with specially designed hardware. Reed-Solomon is used primarily because it is the best way to cope with impulse noise, but it also improves performance in gaussian (continuous) noise — boosting the noise immunity by about 3dB. Trellis coding on its own gives a similar improvement, but together they provide about 4dB.

This coding process is repeated for up to 249 active sub-carriers. The sine and cosine levels for the first five and the last one of the 256 possible sub-carriers are set to 0 and sub-carrier 64's levels are set to predetermined levels so it can be used to synchronise the receiver.

The DSP processor performs a Fast Fourier Transform (FFT) on these to produce a set of 512 14-bit digital values which describe a waveform which is the sum of all the sine and cosine waves of the sub-carriers. This waveform constitutes one symbol, and it may encode more than 2,000 bits of data.

In principle, this waveform will be transmitted to the twisted pair during the next 1/4,000 second while the next waveform is being calculated. In practice, there is a need for a small time gap between one symbol and the next, so there is a 14.5 microsecond 'cyclic prefix' transmitted before the main waveform. This consists of the last 32 samples of the waveform, prepended to the start, so that 544 samples are sent in each 1/4,000 second. The cyclic prefix is a time when echoes (in the twisted pair) from the previous symbol's waveform are allowed to diminish somewhat before the receiver starts sampling the new symbol's waveform. This would give a sample rate based on 4,000 symbols per second and 544 samples per symbol. There needs to be a special synchronisation symbol, which carries no data, to maintain sync after disruptive noise. One distinctive synchronisation symbol is inserted every 68 normal symbols, so although the data frame rate is 4,000 per second, symbols are sent at 4,058.8 per second.

Each sample of each symbol is sent to a 14-bit DAC (Digital-Analogue Converter) at the sampling rate of 2.208 MHz to produce the analogue downstream signal. The voltage from the DAC is fed through a simple filter and then to the hybrid and POTS splitter which combines it with the POTS signal from the analogue phone exchange (*hybrid* is the name given to the analogue circuit in a modem or phone which separates upstream and downstream signals). The hybrid and POTS splitter also separates the received downstream signal from the twisted pair loop in a way that minimises crosstalk from the upstream signal it is driving into the loop.

The peak signal level of a multicarrier signal is much higher than its average level — just occasionally lots of the sine and cosine waves add up in the same direction. The draft standard allows the signal to clip at the limits of the DAC's range on average once in every 100,000 samples. This distortion degrades the transmitted signal and

Figure 1: Echo Cancellation

With echo cancellation, the lowest frequencies are used for both upstream and downstream data. The number of bits which can be encoded on those sub-carriers which share the same frequencies is limited for two reasons: NEXT from similar signals on other cable pairs, which adds to the noise level in these frequencies; and - because it may not be possible to completely cancel the echoes of the upstream signal - these remaining echoes also add to the interference from NEXT. Each sub-carrier is modulated with QAM and can carry between 2 and 12 bits of data per symbol at 4,000 symbols per second. Sub-carriers not used for data or synchronisation are turned off. In the example below, the upstream rate is 488Kbps and the downstream rate is 5.416Mbps.

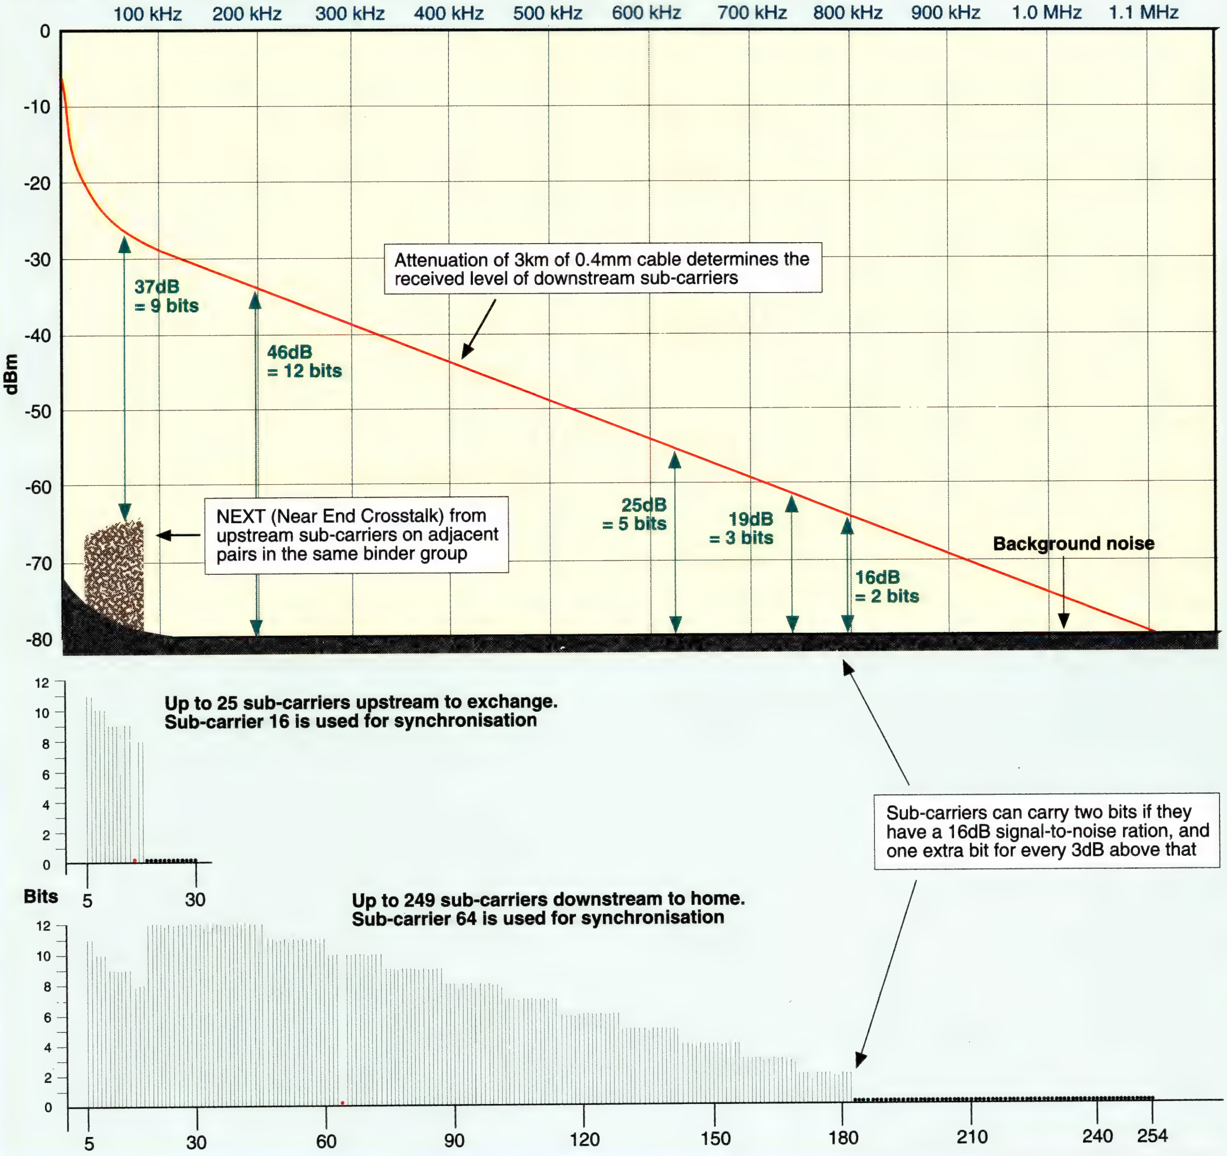

Attenuation of 3km of 0.4mm cable determines the received level of downstream sub-carriers

NEXT (Near End Crosstalk) from upstream sub-carriers on adjacent pairs in the same binder group

Up to 25 sub-carriers upstream to exchange. Sub-carrier 16 is used for synchronisation

Sub-carriers can carry two bits if they have a 16dB signal-to-noise ration, and one extra bit for every 3dB above that

Up to 249 sub-carriers downstream to home. Sub-carrier 64 is used for synchronisation

for the receiver this distortion is indistinguishable from impulse noise spikes at the upper and lower extremes of the symbol's waveform. Since the distortion impulse is short, it has a broad spectrum and its energy is spread thinly over the sine and cosine levels of many of the sub-carriers. Any bit errors it creates are likely to be minimal and so can be corrected by Reed-Solomon Forward Error Correction.

The signal travels down the twisted pair, suffering noise, attenuation, echoes and crosstalk. At the home, it enters a POTS splitter which separates it from the POTS signal and sends it to the ADSL transceiver. The POTS splitters at each end also add to what the receiver sees as the loop's impulse response. In the ADSL transceiver's analogue front end, the downstream signal passes through a hybrid circuit which attempts to isolate it from the upstream signal it is driving onto the line.

A hybrid cannot remove all the transmitted signal because some of it is reflected back from the loop itself — particularly the far end of the loop, bridge taps and changes in cable type. These echoes are tolerable in a phone, or a low speed modem, but a DMT transceiver must remove them because the echoes mix the signal from one symbol into another.

The received signal is typically weaker than the echoes which are mixed up with it and although there are digital methods to cancel them, these echoes make it difficult to digitise the received signal in the first

Inside DMT (Cont.)

Figure 2: Frequency Division Multiplexing

With Frequency Division Multiplexing the sub-carriers used for upstream and downstream data use separate frequencies. Due to the limitations of practical analogue filters, there must be a `no man's land' which cannot be used by either signal. In this imaginary example, the upstream rate is 528Kbps and the downstream 4.560Mbps.

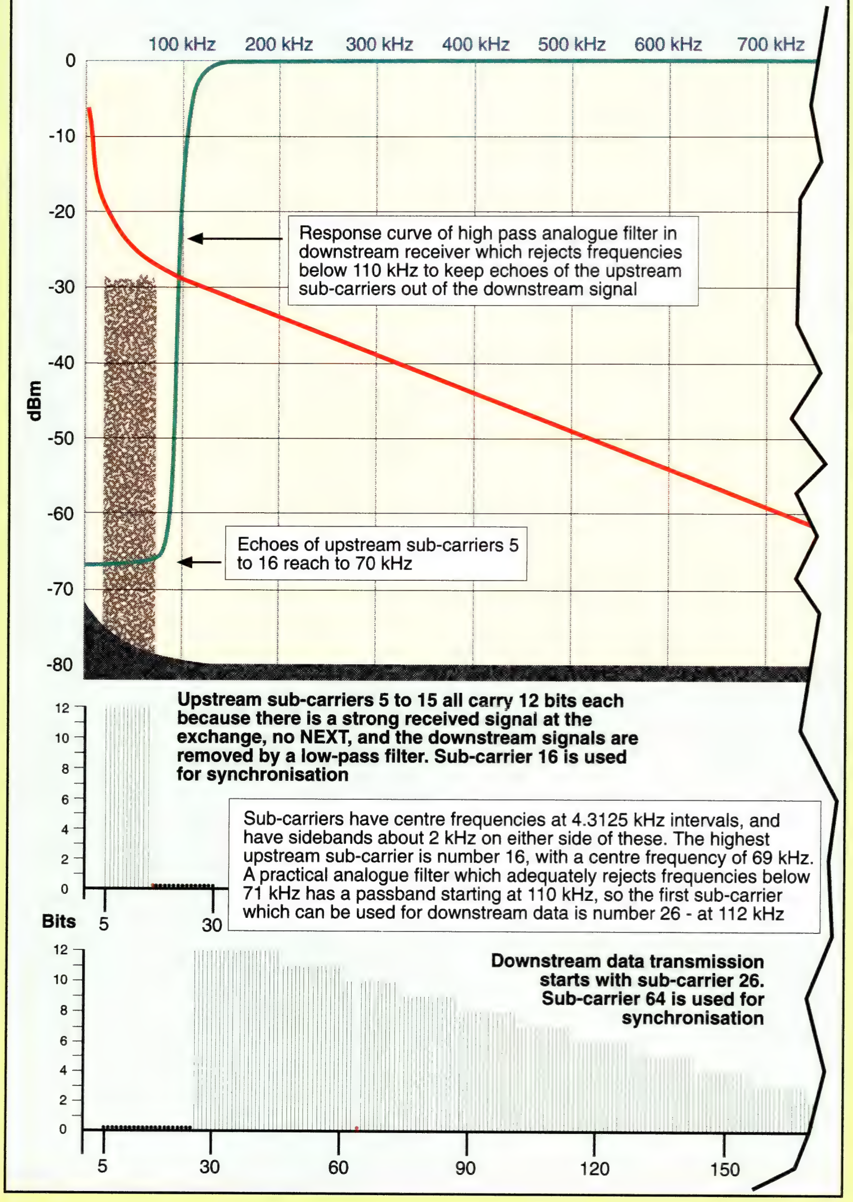

place. The received signal must be digitised in great detail — with a 14-bit ADC (Analogue-to-Digital Converter). If the echoes were three times bigger than the received signal, then an ADC with a four times greater range would be required — a 16-bit ADC, which is prohibitively expensive.

One approach to removing the echoes is echo cancellation (see Figure 1). The downstream receiver maintains a mathematical model of the loop's impulse response — the pattern of echoes it generates from signal sent upstream to it. This model is derived from measurements made during training, but is constantly updated during operation to take account of changes in loop behaviour. The digital representation of the upstream signal is converted into a voltage with a DAC (Digital to Analogue Converter) and sent to the hybrid and so to the loop.

This digital representation is also processed by the model of the impulse response and the results are sent to a second DAC to produce an echo cancellation voltage. Ideally this voltage is exactly the same as the echoes which are being received, and by subtracting this voltage from the received signal coming from the hybrid, the echoes are completely removed.

In ADSL, the lowest sub-carriers are a precious resource. They have the highest signal-to-noise ratios and so can carry the most bits. Echo cancellation provides the highest performance because upstream and downstream signals can both use these lowest frequencies.

FDM (Frequency Division Multiplexing) is easier to understand and design (Figure 2). The upstream sub-carriers use the lowest frequencies and the downstream signals start somewhere above that. In the downstream receiver in the home, this enables echoes to be kept out of the received signal sent to the ADC by using an analogue high pass filter to reject all the frequencies used by upstream sub-carriers. In the upstream receiver, a low pass filter is used to reject the echoes of the downstream signals.

Depending on the upstream data rate and the signal to noise ratio, the transceivers could use up to sub-carrier 30 for upstream data, but at a minimum, they will use up to sub-carrier 15, because sub-carrier 16 must always be turned on to maintain synchronisation.

If it was practical to use analogue filters with an extremely sharp cut-off, then it would be possible to start the downstream data on the next sub-carrier above the highest one used for upstream. However for reasons of cost, noise and distortion, this cannot be done — the best practical analogue filter has a cut-off curve which spans several sub-carriers. These cannot be used for either upstream or downstream data.

Even if echo-cancellation is achieved perfectly, there may be some degradation of the data carrying capacity of the lowest sub-carriers in both directions. This is due to NEXT. The lowest frequency downstream sub-carriers to one home suffer NEXT from the upstream signals from similar transceivers nearby which drive adjacent pairs in the cable buried in the street. A similar situation happens at the exchange. The loss of bits on these sub-carriers depends on how many disturbing pairs there are and how weak the received signal is.

Despite NEXT, echo cancellation is always likely to provide more bits for upstream and downstream than FDM under similar conditions.

Echo cancellation is technically difficult and the Aware-Analog Devices partnership believes that it is too costly to implement. This position may make sense if its higher performance is not required — as it may not be on the majority of loops. Amati's John Cioffi, however, says that a properly designed echo cancellation scheme will be cheaper than FDM as well as providing higher performance.

The signal is passed through a relatively simple analogue anti-alias filter to a 14-bit ADC sampling at 2.208 MHz. In the future, '1-bit' sigma-delta ADCs may achieve similar performance by sampling at 32 or 64 times this rate. By one means or another, the received signal is digitised and the echoes of the upstream signal are removed.

Now that the received downstream signal has been digitised, it is fed through a digital filter which compensates for that part of the loop's downstream impulse response which is longer than the 14.5 microsecond cyclic prefix. Amati's John Cioffi says that this patented technique is essential to achieving DMT's near optimal performance. He won't explain the exact form of the filter but he says it is not excessively complex to implement.

The result is the downstream signal, substantially free of the effects of the loop's downstream impulse response and the echoes of the upstream signal. For the purposes of data recovery, the sync symbol and the cyclic prefixes are discarded. This leaves symbols — each represented by a 512 sample waveform. These are fed to an Inverse Fast Fourier Transform which produces the levels for the sine and cosine waveforms of the active sub-carriers. These values are fed to a Viterbi decoder which knows the behaviour of the transmitter's encoding scheme — QAM or trellis QAM. The decoder recovers the data bits as best it can.

The received bits — which contain some errors due to continuous and impulse noise are split into fast and interleaved buffers. After de-interleaving, the two buffers are processed by the Reed-Solomon Forward Error Correction algorithm. The results are descrambled to produce the fast and interleaved data bytes and their CRC bytes. The CRCs, together with the results of the Reed-Solomon algorithm, are used to check that the data has been recovered successfully. The management data is sent to the local controlling software and the remainder — the downstream simplex and duplex user data — is sent to user devices. These might be an MPEG-2 video and sound decoder system, an ISDN S-Bus interface chip or a personal computer interface — or an ATM switch which directs packets to these devices.

Robin Whittle

From page 89

They acknowledge that this is an ambitious plan for a huge chip.

DMT's extraordinary performance is dependent on a large amount of Digital Signal Processing. This is too much for even the fastest standard DSP processors which execute 40 million operations per second and cost over $US100 each — so several must be used together. However these DSP functions could be implemented more elegantly on a large purpose-designed chip which could be mass produced at a lower cost.

The analogue front end of DMT transceiver is also extremely challenging. This includes the POTS splitter, which uses discrete inductors and capacitors, and three functions which can be integrated onto silicon. The first is the hybrid — which works with the POTS splitter to transmit signals onto the pair, whilst picking up received signals with a minimum of crosstalk from the transmitted signal. The second is the DAC (Digital-to-Analogue Converter) which converts the transmit signal from numbers into voltages. The third and most demanding function is the ADC (Analogue-to-Digital Converter). Associated with all these are filters and/ or analogue echo cancellation circuits which can be quite complex.

The ADC is the most exotic part of the DMT transceiver. In the home transceiver, it must sample at 2.2 MHz and typically needs 14 bits of resolution. The only such devices available now are $US280 multi-chip modules which consume 3 watts. Neither Amati nor Motorola want to discuss the ADC or their techniques for echo-cancellation.

There are at least two other DMT developers. One is Alcatel in Belgium and the other is a partnership between Aware and Analog Devices in the United States. Aware is providing the DSP software and Analog Devices the hardware for a DMT transceiver which they hope to release in August. It will use an unspecified number of Analog Devices 21062 DSPs — a new high performance 32-bit floating point CPU which is designed to exchange data rapidly with other 21062s.

Analog Devices will not be selling transceiver boards, but will sell the chips, schematics and the software on the open market. Three companies — AT&T, Westell and Newbridge — are first in line and will make standalone transceivers and sell them to the US RBOCs and other telephone companies around the world who are desperate to trial DMT. This initial transceiver will be too expensive for domestic sales, but will enable trials to take place whilst the software is improved and the standard evolves.

Marcos Tzannes, a software engineer at Aware, says the company has decided against using trellis encoding and echo cancellation, preferring to concentrate its efforts on enhancing other aspects of the transceiver — particularly the analogue section. Analog Devices' Marketing Manager, Greg Whelan says "All of the operating companies want low low cost and not fancy features."

Trellis encoding and echo cancellation extend the reach of the system and so may not be needed for the majority of customers. They are difficult to achieve, but John Cioffi says his company Amati can include them in products at little extra cost. He says Amati has a proprietary echo cancelling design which requires only a 12-bit ADC, which is a lot cheaper than a 14-bit ADC.

Analog Devices is a leader in ADC and DAC technology. Dave Robertson, an engineer on the ADSL project, was happy to talk about Analog Devices' approach to the ADC problem. It has developed a single chip 14-bit 2.5 MHz ADC which consumes 600 mW. This ADC is a quarter the power of the 14-bit multi-chip module ADCs available today. The chip has been fabricated and tested and will sell on the open market next year for around $US70 and 'a lot less' as part of an ADSL chipset.

DMT ADSL needs a relatively cheap ADC such as this before it can become cost-effective.

However Analog Devices will not rush to build its converter and DSP technology into a single chip transceiver. It plans a multi-generation approach, perhaps with lower cost 2Mbps systems. Robertson says: "Our approach will be to have the most economic solution whether that is just one chip or four chips. There is nothing magical about a single chip, what's magical is the price." He offered no estimate of when the development of a dedicated DMT DSP chip might begin, but chips of that size typically take at least a year to get into production.

The price will be magical only if it is low enough, soon enough. It seems possible that DMT chipsets — or chips if Motorola is successful — may start appearing in 1996. It is impossible to predict what their prices will be, but $US150 seems a reasonable estimate.

Amati plans two new products. Later this year it plans to produce the Presto, a 2.048Mbps DMT transceiver with a much lower cost and a longer range than 2Mbps CAP systems. Prototypes of the second product — the Overture — should be available by the end of 1994. It will use VLSI chips which Amati has been designing since last year and will use trellis coding and echo cancellation. It will go beyond the current draft standard and provide up to 9Mbps with upstream data rates as high as 800-Kbps. These higher rates anticipate the addressing overhead inherent when using ATM to manage the data streams.

There is no shortage of expertise or funds in the race for a cost-effective DMT transceiver. However claims for DMT's early cost effectiveness should be considered in

MPEG-2 Compression: 6Mbps Required For Quality

The MPEG-2 (Moving Pictures Expert Group) standard specifies a digital language for video compression based on a number of techniques for building frames directly, or by reference to other frames.

The compression techniques used in MPEG-2 are complex to say the least. In simple terms, the encoder instructs the decoder how to build an initial 'I' frame from scratch using techniques based on an algorithm called the Discrete Cosine Transform. Colour and luminance are treated separately. It can tell the decoder how to modify the data in an 'I' frame to make a new 'P' frame — typically by moving 16 x 16 pixel blocks of the 'I' image around and then colouring in some more details. Similarly this 'P' frame can be modified to produce a subsequent 'P' frame.

A third type of frame — a 'B' frame — can also be specified. This frame is based on interpolating between two 'I' or 'P' frames and adding more details. These techniques are used in both the MPEG-1, and MPEG-2 standards. (In last month's article it was incorrectly stated that MPEG-1 did not use 'B' frames.) 'I' frames are typically used when the image changes drastically. Even if there are no major changes from one frame to the next, an 'I' frame is sent about 5 times a second so that errors in the received data do not propagate for more than a fraction of a second.

For a given program, the quality of the encoded images depends primarily on how many bits per second can be used. The encoder has to decide which of the many details in an image to concentrate most bits on, so another factor is how 'smart' the encoder is. For example, if a compere was sitting in an audience, the viewer would want more bits dedicated to his face than to the faces in the crowd, but a real-time encoder has no way of knowing this.

For a given bit rate, the quality produced by a real-time encoder is constrained primarily by its computational power and its limited knowledge of what the viewer wants to see. Non real-time encoding can be done at a more leisurely pace, so there is no hard limit on the computational power available. In addition, a human operator can intervene to direct the encoding algorithm to concentrate more detail where the viewer is most likely to want it — for instance on a person's eyes and lips rather than on the moving background.

Subjective Viewing Tests

The question of data rates for acceptable MPEG-2 encoded video can be answered only approximately. There are many technical and human factors and the program material varies greatly in how easily it can be compressed.

Material sourced from a video camera is the most demanding. There are 25 frames per second, each composed of two fields. The first field includes the odd numbered lines and the second the even. Each field carries image information about events 1/50 second after the previous field — long enough for a rapidly moving object to move appreciably. So it is not feasible to combine the fields into a single image — they must be considered as separate images where there is rapid motion. In contrast to MPEG-1, MPEG-2 enables the encoder to decide whether to treat parts of the image as separate fields or as frames.

However if the source is film, then both the odd and even fields typically come from a single film frame, so they can be combined and considered as one of 25 frames per second. This enables movies to be compressed more easily at a given bit rate.

Both technical and human factors affect the perceived image quality for a particular program at a given data rate. At low bit rates, one encoding approach may produce one set of compression error artifacts and another may result in something quite different. One viewer may notice one kind of error more than another.

These problems are different from the noise, smearing and ghosting which can very often result from analogue transmission and VCRs. They include lack of fine detail, fuzzy areas, fast moving objects disappearing and straight lines developing jagged edges.

For instance, the numbers of a digital clock superimposed on the crowd at a well-attended swimming event were surrounded by a fizzy, multicoloured crosshatched halo. In another case, a newsreader's face and hair is reproduced in remarkable detail as she moves against an even background, but the freckle on her neck moves a third of a second after she does! These problems are visible at 1.7Mbps. But at higher data rates — such as 6 to 9Mbps — they would disappear.

Another problem arises when material is subject to more than one stage of compression. Perhaps it is compressed and decompressed on its way to the studio, and then compressed again — not necessarily with the same algorithm on its way to the home. The artifacts from the first compression add extra unwanted detail to the input of the second — which may distract it from concentrating on the details the viewer wants to see.

Telecom Research Laboratories is contributing to the ISO/IEC MPEG-2 standardisation effort and has developed compression and decompression algorithms which are implemented in software and run on a high-speed workstation at a frame every few seconds. The results are transferred to digital video tape for viewing. Files are transmitted over the Internet to other companies doing similar work so that each can verify their algorithm's compliance with the standard and optimise their performance.

In Canberra, the Communications Lab of the Department of Communications and the Arts is a member of the MPEG-2 test group. In February it conducted viewing tests according to ITU-R recommendations using 625-line 50 Hz video compressed at 9, 6 and 4Mbps. No sound was used in these tests. Four other organisations overseas conducted tests on the same material and also investigated 525-line 60 Hz video at 9, 6 and 4Mbps and at 2Mbps for material sourced from film.

The results from these tests were considered by the MPEG-2 test group in March and one of the group's conclusions was that 'Services having a quality comparable to that of present PAL require at least 5 to 6Mbps at 625-line/50 Hz. Services operating at a low bit rate (4 or 2-Mbps) do not have the same potential quality as normal television.'

For a given data rate and compression algorithm, the subjective viewing results depend on the type of program material, the viewing distance and how critical the as-

the light of [Stewart] Fist's Law: *"A product takes twice as long to develop as planned. When it arrives, it costs twice as much as first claimed — and is half as good."*

Total System Cost

Two transceivers are required for each ADSL service — one in the home and one at the exchange. However, the one at the exchange has less receive processing to do and its ADC only needs to sample at 276 kHz.

In the home, the set-top will need a complete MPEG-2 decoder chipset. David Taylor, Marketing Manager for C-Cube Microsystems, says that MPEG-2 decoder chips will go into startup volume production in the second half of this year, with a price of $US50 to $US60. He expects the price to fall to $US30 to $US35 as the market expands to about a million chips a month in 1995. The decoder needs a PAL driver chip and four 4 megabit DRAM chips — which may cost $US10 each or so in 1996.

Crystal Semiconductor has a DSP DAC combination chip in production already which is optimised for MPEG-2 stereo audio decoding. Its cost now is $US38. So a sessor is. Some types of program material are significantly more affected than others. In the 625/50 tests, for example, a demanding basketball sequence scored around twice the average degradation of other material in the group's tests. This was much more of a problem at 4Mbps than at 6Mbps.

Film is easier to compress than video at a given data rate, but there is some problem converting its 24 frames per second to our video's 25 per second. The problems are significantly worse for 30 frames per second used in the US and Japan.

TRL's Dr Michael Bigger has some video sequences compressed at 1.7Mbps, which may not be comparable to the material used in the subjective viewing tests. Viewing these as part of researching this article I found the granularity and shakiness of detail were occasionally obvious, and always visible if I looked for them.

In the basketball sequence, I kept my eye on the ball and didn't notice any problems — which surprised me. I thought few people would notice the loss of detail in the parquetry floor, but other viewers may concentrate on details I missed, such as facial expressions. Viewers may find the artifacts of low bit rate compression — such as shimmering — to be distracting and annoying.

TRL has not yet investigated MPEG-2 audio compression, but according to Dr Biggar "the consensus is that 256 or 384-Kbps is fine for 'almost CD quality' stereo sound." This is comparable to the compression ratio of the Sony MiniDisc — which does 6 to 1 compression of the 1.411Mbps CD rate. It may not pay to skimp on sound quality — several independent studies have shown that viewers perceive higher image quality when shown the same video with improved sound quality.

So while movies and educational material may be acceptable to viewers at 2 or 3Mbps, it is clear that 4 or 6Mbps will be essential to achieve acceptable quality with other types of program — particularly live sport. However these are early days for video compression — the real test will come with day-to-day programs and paying customers.

Managing Differing Rates

Viewers may be connected to a Video Server via ADSL at 2, 4 or 6Mbps, or by digital coax or fibre at these and other rates. They may wish to view one or more programs via this channel, so ideally the server should be able to provide a given program at a variety of data rates.

One approach is to store a separate file for each data rate — say 2, 3 and 6Mbps, with consequent increase in storage requirements. A second approach is to store a limited quality file at 2Mbps and a supplementary file at 2 or 4Mbps which adds finer details to the basic image. TRL was one of five organisations which is working in this scalable approach, but it is not part of the 'main profile' of MPEG-2 which is being developed by 20 organisations. The scalable approach involves significant extra costs at the encoding, storage and network management levels, but the extra costs at the decoder are minimal.

Asynchronous Transfer Mode specifies a packetising and switching protocol based on 48-byte data packets with relatively short 5-byte address headers. These short fixed length packets can be switched directly by hardware, rather than by storing the packet and using software to determine where it should be routed to.

In March, Telecom awarded a $2.5 million contact to Australian ISDN specialist Jtec to develop ATM ADSL switching technology. As this issue went to press, it was still too early for Jtec Managing Director John Riedl to discuss the details of the project. However, he says that ATM can solve the problem of managing variable data rates and is more network efficient — particularly with interactive services. Interactive applications involve bursts of data to update a screen, followed by seconds or minutes of inactivity.

ATM provides a number of architectural options for the Video on Demand and interactive services network. It may influence how the material is encoded and stored, how it is transferred through fibre, coax and ADSL, and how it is handled at the home where it may be addressed to a number of logical devices.

Robin Whittle

---

complete MPEG-2 chipset with supporting components may cost as little as $US110 in late 1996.

To calculate the cost of the home electronics we need to add the costs of the MPEG-2 and DMT chipsets — $US110 plus $US150 — to the costs of the supervisory microprocessor, remote control, power supply, and supporting mechanical and electronic hardware. These include the POTS splitter in a separate box, the infra-red remote control and cabling. Then we add the costs of manufacturing, test and after sales support and factor in R&D, profit, markups and sales tax. If the retail price approaches $1,000, then it's back to the drawing board!

It is hard to see how there could be much change from $1,000 in the first few years of production. Current 14.4Kbps modems are based on a chipset costing $US100 or less. They are a mature, mass produced product with perhaps a tenth the complexity of an ADSL set-top, but they retail for around $500.

An ADSL service to each home will involve well over $1,000 of investment by the customer and the telco to cover the cost of the set-tops and the transceivers and infrastructure at the exchange. For VOD, the share of the Video Server costs will be significant too. However, when VOD is available, rather than just Pay TV, consumer demand could be very strong indeed.

The ADSL component of this may well compare favourably to coaxial cable where coax is difficult to install. Coax requires significant investment just to run it past homes, so it could be expensive unless the uptake rate is high. However, ADSL can reach the early adopters scattered in the suburbs without the up-front expense of laying new cable.

Let's Not be Hasty

It seems likely that DMT ADSL will be able to provide 6Mbps to between 80 and 95% of homes in Australia's cities and towns — and 2Mbps to most of the remainder. This will require a major investment starting around 1997 — which is how long it will take for the technology to mature. Perhaps by then video servers will be cost effective and Video On Demand will drive ADSL and digital coax adoption.

It can reach most of the early adopters without the need to lay cable throughout entire suburbs. It can provide Pay TV, VOD and interactive services to the growing number of people who will want them — but not to everyone, and not at rates fast enough for HDTV. For the next 15 years or so, both digital coax and ADSL may connect us to the Information Super-you-know-what and by then there will be irresistible reasons to run fibre to every home.

In the meantime, we should be cautious of computer models and the always ambitious claims of developers. We should also be wary of pioneering attempts to deploy immature technology — which could prove so successful that we are locked into a substandard service for a decade or two. For both digital coax and ADSL, the target is at least 6Mbps downstream, with a back channel of at least 400Kbps — for video conferencing — after allowing for the overhead of ATM which is clearly the best way to manage these resources.

Robin Whittle is a Melbourne-based electronic designer and software programmer. The author wishes to thank Necdet Ozergun from Telecom Research Laboratories, John Bingham and John Cioffi from Amati and Michael Johannessen from NEC for their assistance in preparing this article.

300% faster than Cisco.
50% faster than Olicom.
60% faster than IBM.
50% of the Cost.

Introducing the ACSYS TokenSwitch 3000 Bridge/Router

"It's the fastest Token Ring bridge we've ever tested." Kevin Tolly, The Tolly Group

The numbers are in. The highest performance Token Ring bridge/router is the new TokenSwitch 3000 from ACSYS. It's the only bridge that is wire speed at 64 byte packets and up – at a per packet cost that's half the competition's. Highest throughput rates and ultra-low latency make it the ideal choice for client/server applications requiring fast response time...and for reducing dropped sessions if you have integrated SNA onto the LAN. TokenSwitch 3000 is IBM® compatible, supports both LAN Network Manager™ and SNMP management. It combines the sophistication of IPX™ and TCP/IP routing with the speed and simplicity of source/route bridging.

Get your free copy of The Tolly Group benchmark and our free **Token Ring Performance** white paper.

Call ACSYS today.
You'll be 100% impressed.

Any Questions?
Call 1-800-GO-ACSYS

Singapore
Ph Int 65 432 0311
Fx Int 65 480 4535

Australia
Ph Int 61 2 299 2799
Fx Int 61 2 299 2807
Toll Free 008 63 3337

New Zealand
Ph Int 64 9 634 4580
Fx Int 64 9 634 4586

ACSYS
An AMP Company

Kevin Tolly and
David Newman

Testing Remote Token Ring Bridges

Remote access may be all the rage, but with the high cost of wide area network bandwidth, it's no surprise that network managers want remote bridges that can make the most of links between sites.

Here's a little quiz: what are the two most expensive features of any remote Token Ring bridge? The RISC processors? The redundant ports? The flashy LED displays? Guess again. The high cost of WAN bandwidth and the often expensive administrative charges associated with remote connectivity are the real big ticket items. Small wonder, then, that network managers want remote bridges that make the most of bandwidth and don't need an army of tech support engineers to keep in service. The management issues are simple enough to resolve: buyers are demanding easy-to-use, plug-and-play products that can be remotely managed from a central site.

And what about those two-tonne phone bills? Luckily, data compression has emerged as the key way to avoid expensive upgrades to faster WAN lines. According to vendors, compression makes it possible to squeeze text files by a 4:1 ratio — pretty much the stock claim industry-wide. This can result in big savings for corporations, since compression can mean fewer WAN circuits. Compression is so important, in fact, that in this evaluation of seven remote Token Ring bridges, the Data Comm Test Lab only considered products that include the technology.

The Big Squeeze

The Lab wanted to see for itself just how well vendor claims matched up to reality. What we found was that only one remote Token Ring bridge — the ACS 4200 from Advanced Computer Communications (ACC) — actually attained 4:1 compression with the text files used by the Lab, earning it a Tester's Choice award. Three other bridges compressed text at better than 3:1, including the Remote Bridge 200 from Develcon Electronics, which also garnered a Tester's Choice for its combination of ease of use, good performance, and compact size. Of the remaining three bridges, none could top 2:1 compression.

Because compression ratios depend largely on the data involved, buyers of the bridges we tested may get better or worse results than those reported here. Since we tested all bridges under the same conditions, however, these tests should offer prospective buyers a fair way to compare performance. Judging from the wide range of compression ratios we found, it clearly pays to shop around.

But performance alone isn't the whole story. The decision to buy a bridge is based on many factors, including the number of WAN ports, the maximum link speed, and the highest speed at which compression is possible (see Table 1). Whether compression is implemented in hardware or software can be important at higher bandwidths. Filtering, which helps keep internetworks from getting overloaded, also should be evaluated. And, as always, price needs to be taken into consideration.

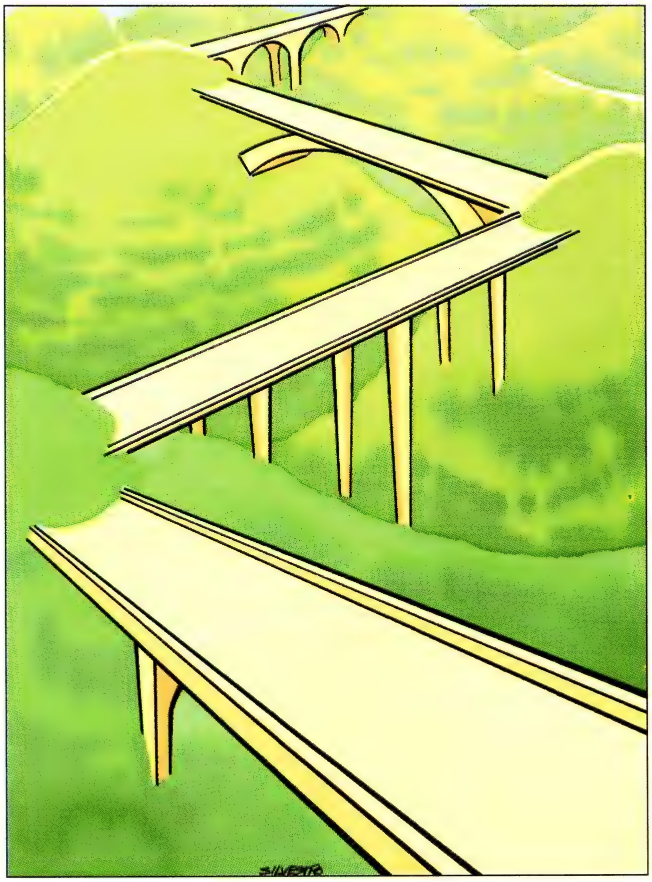

To help network managers put all these factors in perspective, we assessed each bridge according to its ability to perform several critical tasks and rated it on a scale of 0 to 10 (see Table 2 on page 100). We tried, for example, to set a filter that blocked traffic from one MAC address. If we were able to do so without referring to

Table 1: Selected Vendors of Remote Token Ring Bridges

VENDOR	PRODUCT	MAXIMUM LAN/WAN PORTS†	WAN SPEEDS	MAXIMUM COMPRESSION SPEED	COMPRESSION	FILTERS
Acsys (02) 299 2799	Tokenmaster 5000	1/4	9.6Kbps to 2.048Mbps	1.544Mbps	Hardware	MAC, SAP, NetBIOS, Link Limiting, SNAP
Advanced Computer Communications ADE Network Technology (03) 543 2677	ACS 4200	1/2	1.2Kbps to 2.048Mbps	256Kbps	Hardware	MAC, SAP, NetBIOS, SNAP
Andrew Corporation (03) 357 9111	Pathwise/7622	1/2	9.6Kbps to 2.048Mbps	256Kbps	Software	MAC, SAP, NetBIOS, SNAP, custom
Develcon Electronics Anderson Network Corporation (02) 838 7755	Token Ring Remote Bridge Model 200	1/2	9.6 to 512Kbps	256Kbps	Hardware	MAC, SAP, NetBIOS, SNAP, custom
Microcom Systems (02) 410 9953	MBR/6003	1/2	9.6Kbps to 2.048Mbps	1Mbps	Hardware	MAC, SAP, custom
Olicom Force Technology (02) 971 1000	Token Ring Remote Bridge 16/4	1/1	9.6Kbps to 2.048Mbps	128Kbps	Software	MAC, SAP, NetBIOS, Link Limiting, custom
3Com (02) 959 3020	Netbuilder Token Ring Remote	2/2	1.2Kbps to 9.216Mbps	64Kbps	Software	MAC, SAP, NetBIOS, SNAP, custom

LNM = LAN Network Manager
SNAP = Subnetwork Access Protocol
MAC = Media Access Control
SAP = Service Access Point
† Maximum for model tested; more ports are available from some vendors

documentation or calling the vendor, we gave the product a 9 or 10, depending on how long it took. If the task involved a bit of navigation through configuration routines, we rated it from 6 to 8, again depending on the time involved. Having to refer to the documentation merited a rating of 4 or 5, and the worst-case scenario — calling the vendor — meant a rating of 3 or less. We also rated the clarity of documentation.

Building Better Bridges

Remote Token Ring bridges aren't plug-and-play devices yet, but they're getting there. Indeed, this group of bridges generally proved to be far simpler to install and configure than the last batch we evaluated.

But there's still room for improvement. Even though our test bed reflected the most common configuration for this gear — two 16Mbps LANs connected over a WAN link — none of the products were ready to go right out of the box. In some cases, the Token Ring data rate had to be changed and at times it was difficult to turn compression on or off. And two bridges didn't allow us to set the Token Ring frame size — a critical flaw that can choke a low-speed WAN link and force session timeouts.

Vendors take several approaches to configuration. The simplest schemes use front-panel switches or Windows-based configuration programs. Acsys offers both with its Tokenmaster 5000. The bridge/router's front-panel buttons and LCDs allow net managers to set virtually all LAN and WAN parameters. And the vendor's Netacsys configuration program is clearly and logically organised, helping to speed and simplify the whole process.

But not every bridge/router is as easy to configure. The Netbuilder Token Ring from 3Com is a good case in point. The product supports the same extensive list of network-layer protocols and other features as the vendor's high-end router, the Netbuilder II. But sorting through all the options can prove perplexing, especially when the device is being used only for bridging. To be fair, we should point out that 3Com's configuration utility is far easier to use than those from competitors like Proteon or Wellfleet Communications.

Rating Remote Bridges

To get a real feeling for configuration, we chose key LAN and WAN settings and graded them on how easy they were to alter. We started on the LAN side, where we attempted to change ring numbers and MAC addresses — something network managers would have to do to accommodate network growth or replace a failed interface card.

On the wide area network side, we attempted to change three key parameters: maximum frame size, speed of the WAN link, and status of data compression. Frame size is critical on all wide area network links, especially on low-speed ones. Larger frames allow greater throughput, but one big frame can block a WAN link long enough to cause a session timeout. Obviously, network managers need a way to control the frame sizes.

Two of the bridges in this test — ACC's ACS 4200 and Microcom's MBR/6003 — offer no way to change WAN frame size. Microcom says its bridge can handle frames up to Token Ring's theoretical limit (approximately 17,000 bytes), but net managers have to set sizes on the LAN side of the bridge. That's not always possible, however. The source routing protocol requires that the *bridge* control the maximum frame size. Many applications send the largest frame size allowed, and then expect the bridge to throttle back the size if necessary. If the bridge can't be set, there's no way to control the size of frames sent onto the wide area. In effect, the approach taken by Microcom is actually a protocol violation.

Clocking — the rate at which data is sent onto the WAN link — isn't usually handled at the bridge. Instead, that task often falls to attached network termination units. But some bridges automatically sense the speed of the WAN link and set their WAN ports accordingly. Bridges with this capability are available from ACC, Acsys, Develcon, and Microcom. The other bridges required the WAN link speed to be defined, regardless of where clocking is determined. The lack of dynamic clocking can be a problem especially for bridges with dial backup capabilities. If the primary WAN link fails, the bridge will have to be reconfigured before it can establish the dial-up connection — obviously an unacceptable state of affairs.

Since most products come with compression enabled as a default setting, we wanted to get a sense for how tough it is to

SUPPORT FOR 10 INDIVIDUAL/10 GROUP FILTERS	NETWORK MANAGEMENT
Yes/Yes	LNM, SNMP
Yes/No	SNMP
Yes/Yes	LNM, SNMP
Yes/Yes	LNM
Yes/No	SNMP, HP OpenView
Yes/Yes	LNM, SNMP
Yes/Yes	SNMP

disable this feature. For Acsys and Develcon, the task is accomplished with a few keystrokes. But users of 3Com bridges can expect to wade through several levels of menus. To be fair, it's unlikely that most net managers will want to disable compression. In some cases, though, poor line quality may be a problem when the compression is on; in others, the higher effective rates allowed by compression may be more than attached devices can handle.

No Through Trucking

Data compression may be a great way to squeeze more data into the WAN pipe, but there's some traffic that doesn't belong in the wide area in the first place. Depending on the applications in use, overhead traffic like source route broadcasts and SAP (service access point) frames can consume significant amounts of WAN bandwidth. Bridge filters make it possible to restrict traffic based on a variety of preset or user-defined conditions.

The simplest filters block traffic based on source or destination MAC address, or both. A network manager may want to filter out traffic from stations running an application that requires lots of acknowledgments. A few bridges' MAC filters also work the opposite way, passing only those MAC addresses the network manager specifies.

Some bridges also can filter groups of MAC addresses. This allows net managers to set software MAC addresses by ranges. Group filtering is available from Acsys, Develcon, Olicom, and 3Com.

In some instances, network managers only want to restrict some of the traffic from a given station, and there are filters that can block traffic by application or protocol. All the bridges tested except Microcom's can filter frames based on NetBIOS data in each frame's header — a critical way of cutting down on overhead on SNA internetworks. IN addition, the Acsys and Olicom bridges also support IBM's Link Limiting filter, which controls the number of active LLC2 (Logical Link Control 2) sessions on a WAN link. All the bridges in this test can filter the NetWare SAP frames. And many filter the SNAP (subnetwork area protocol) data generated by AppleTalk and TCP/IP stations.

Making it Manageable

Setting parameters on an individual bridge is one thing, but network managers also need a way to monitor and manage installations consisting of dozens or even hundreds of bridges. And when it comes to remote bridges, they also need a way to tell the status of attached WANs and LANs. Fortunately for Token Ring users, they can rely on SNMP and IBM's proprietary (but widely still implemented) LAN Network Manager (LNM) to get the job done. And what is more, Microcom's MBR/6003 also reports to Hewlett-Packard's Openview management platform.

SNMP agents, implemented in all boxes except Develcon's, report bridge status information to any SNMP-compliant network management platform, regardless of the vendor involved. The amount of information sent up to the management system depends on which parts of the bridge MIB (Management information Base) the vendor chooses to implement.

In the Token Ring world, there's a more important net management specification than SNMP: IBM's LAN Network Manager. LNM is significant not only because of IBM's dominance in the Token Ring market but also because of the wealth of data it provides. Among other things, LNM agents report the status of any node on a given ring; monitor remote rings for beaconing or other errors; and test paths across multiple bridges and rings. LNM data also can be passed to NetView, IBM's enterprise-wide management architecture. LNM-compliant bridges that support data compression are available from Acsys, Andrew, Develcon, and Olicom.

Notably, all the bridges tested allow management information to be sent in-hand — that is, network managers don't need to set up another channel to make configuration changes.

Back to Basics

High-level features like configuration and management are very important — but so are seemingly basic physical features like the number and type of ports on a remote bridge.

When we evaluated remote Token Ring bridges back in 1991, some vendors were content to follow IBM's approach of one LAN port and one WAN port (a configuration also known as 1:1 or 'half' bridging). The problem with this scheme is that separate bridges are required at the central site for each remote site. Multiport bridging — also called 1:N bridging — rectifies the situation by allowing one central-site device to communicate with multiple remote sites. These days, multiport capabilities are pretty much a given; all vendors except Olicom supplied 1:N devices.

What's more, bridges from Acsys, Develcon, Microcom, and 3Com balance traffic across redundant WAN links. These vendors say load balancing guards against the possibility of any WAN pipe becoming clogged. Further, most vendors in the test automatically switch traffic from a failed WAN link onto another that's still functioning, And a link needn't stop working altogether for a cutover to occur.

As might be expected, port speeds varied among vendors. On the WAN side, most supported links to E1 (2.048Mbps). Develcon topped out at 512Kbps, while 3Com offered WAN links to 9.2Mbps. On the LAN side, all vendors supported 16Mbps data rates. The Acsys Tokenmaster 5000, however, assumed 4Mbps LAN rates as a default, even when it was attached to a ring operating at 16Mbps.

The Squeeze Play

As a rule, hardware-based compression will outperform software — but not always. 3Com's Netbuilder TR takes the latter approach, yet it posted a second-place finish with a ratio of nearly 4:1. Meanwhile, the Acsys Tokenmaster 5000 compressed data at less than 2:1, despite dedicated compression hardware. What's more, Acsys and Develcon both license hardware compression technology from Stac Electronics, but Develcon compressed data nearly twice as efficiently as Acsys. Clearly, implementations vary widely by vendor, and network managers are well advised to be leery about vendors' compression claims.

To establish a performance baseline, we sent traffic with compression disabled. Throughput for all bridges hovered around 56Kbps.

Then we evaluated compression with two kinds of traffic (see 'Test Methodology' on page 104). To start, we sent ASCII data, which is highly non-random and thus easy to compress. (Our data wasn't totally non-random, though; theoretically, it should be possible to achieve higher compression ratios by using data consisting only of one character, repeated over and over.)

Most vendors say their algorithms deliver compression ratios of 3:1 or 4:1 when

Table 2: Token Ring Report Card

		ACSYS	ACC	ANDREW	DEVELCON	MICROCOM	OLICOM	3COM
Filter Configuration	Set filter to disable 1 MAC address	10	7	8	9	3	5	3
LAN Configuration	Change ring number of local ring	10	8	8	9	7	7	8
	Change MAC address of bridge's local interface	10	8	8	9	N/A	7	N/A
WAN Configuration	Change maximum frame size on WAN side	10	N/A[1]	8	9	N/A[1]	7	7
	Change link speed	10[2]	10[2]	10[2]	10[2]	10[2]	5	5
	Disable compression	10	7	8	9	8	5	1
Support	Clarity of documentation	8	7	7	8	6	5	6

Ratings: Excellent, 9-10; Good, 6-8; Fair, 4-5; Poor, 1-3

1. Cannot be configured on WAN side; frames can be as large as Token Ring allows.
2. Link speed adjusts dynamically in response to external change.

N/A = Not applicable

handling text. As our results show: throughput varied widely by vendor. The top performer was ACC, with a compression of 4.2:1.

Next, we used a precompressed version of the same ASCII data, which, because of its nearly random sequence of characters, cannot be compressed further. Sending this sort of traffic can degrade throughput, since some algorithms may make futile attempts at further compression. The goal of the second test was to check for the presence of an anti-expansion algorithm — a routine that tells the compression algorithm not to attempt to compress uncompressible data. All vendors passed this test; in all cases, rates of about 56Kbps were the maximum for sending precompressed traffic.

Unlike PC-based file compression programs, remote bridges have to perform in real time. That means the bridges' compression algorithm cannot add significant latency (time spend inside the bridge).

Guarding against latency is critical consideration, especially for applications in which every frame sent across the WAN requires an acknowledgment. Theoretically, it's possible for a device to yield very high compression ratios but also to exhibit very high latency, which would be the case for a particularly effective compression algorithm that takes an inordinate amount of time for each frame.

Fortunately, that wasn't the case in this test. When the Lab checked, we discovered a perfect inverse relationship between effective throughput and latency: the lower the latency, the higher the effective throughput. Network managers looking to cut down latency also can switch to windowing applications and increase the window sizes of such applications.

Acsys

Acsys's Tokenmaster 5000 bridge/router is hard to beat for ease of installation, configuration, and management. But compression left a lot to be desired, as evidenced by its fifth-place finish in our throughput tests.

Acsys (formerly Netronix) has put a lot of effort into making its product easy to manage. Its bridge/router can be set up via front-panel buttons and an LCD panel. The latter also shows the unit's vital signs, including port status, WAN link speed, attached ring numbers, and so on. The vendor even labels each port's MAC address. Netacsys, the Windows-based configuration program, is even simpler to use; virtually all configuration changes can be made in less than one minute. Further, the bridge can report status information to IBM's LAN Network Manager and SNMP. Our only quibble with the product's external design was the lack of a protective cover on the front-panel power switch.

We found less to like with the Tokenmaster 5000's internal design. Even though the bridge implements data compression in hardware, the highest effective throughput it could achieve was only 98Kbps — a ratio of 1.75:1. The low rate was all the more surprising considering the Tokenmaster 5000 uses the same compression technique as Develcon's bridge, which achieved virtually twice the effective throughput.

Another problem is that the Tokenmaster cannot be configured from a PC that's also running NetWare; net managers must unload the PC's IPX stack before running Netacsys.

On a more positive note, the Tokenmaster 500 supports up to four WAN ports, the most of any vendor in the test.

ACC

In a field rife with dubious claims about data compression, Advanced Computer Communications actually delivers the goods. ACC's bridge, the ACS 4200, compressed ASCII traffic at a ratio of better than 4:1, the highest of any vendor in this test. And the box combines top-notch performance with straightforward configuration and management utilities. On the downside, the ACS 4200 lacks a couple of key features, which make it somewhat less than 100% IBM-compatible.

ACC's bridge is configured through any ASCII terminal or PC-based VT100 emulator. The character-based configuration routines were relatively simple. The bridge also reports status information to SNMP management systems.

Although the ACS 4200 supports wide area network data rates to E1 (2.048Mbps), data compression can only be used up to 256Kbps. That's not a problem with the majority of leased lines, but it does rule out the possibility of achieving 4:1 compression at high rates. ACC says it is working to raise the maximum link rate for data compression.

Network managers in all-IBM environments will surely notice the lack of support for IBM's LAN Network Manager — an essential net management architecture for Token Ring. The ACS 4200 however, does report to Netview/6000, IBM's SNMP network management offering that runs on its RS/6000 Unix workstation. ACC is planning to add LNM in the third quarter. On a lesser note, the ACS 4200 also does not support IBM's Link Limiting filters, which control the number of active LLC2 sessions on a WAN link.

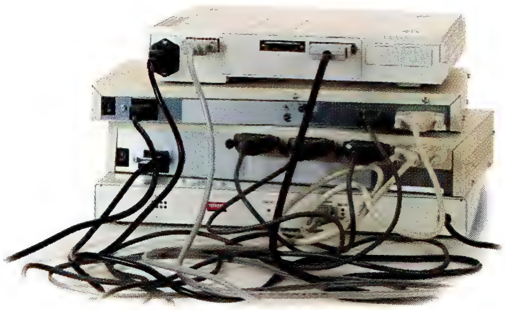

First Stackables.

Wellfleet INSIDE

Now Snappables.

Introducing SnapLAN™, the first remote office hub that anyone can install. And everyone can use.

Stackable hubs made remote LANs possible. Now Fibermux has taken the next step. We made them easy.

Just slide a SnapLAN module into the grooves and listen for the 'snap'. That's SnapLAN's revolutionary internal coupler making all the power and signal connections. Automatically.

No cables. No confusing backpanel wiring. No problems.

The use of flash memory in SnapLAN's integrated Wellfleet router and management module ensures remote operation stays just that. Remote.

Your office staff need never touch the SnapLAN box as all updates, diagnostics and maintenance are done from the central office.

And as SnapLAN can hold any combination of Ethernet or Token Ring module, your remote offices can remain as flexible as they choose.

SnapLAN. Solves remote connectivity and management problems, in a snap.

I need more information on SnapLAN's remote office solution. Please rush me literature on the SnapLAN snappable hub.

Title__________ Name__________ Position__________

Company__________ Address__________

__________ Phone__________ Fax__________

Please fax to: (02) 975 5337

ADC Telecommunications Australia
5B/4 Skyline Place, Frenchs Forest,
NSW 2086 Australia
Tel: (02) 975 1499

CERTIFIED ISO9001 CERTIFIED

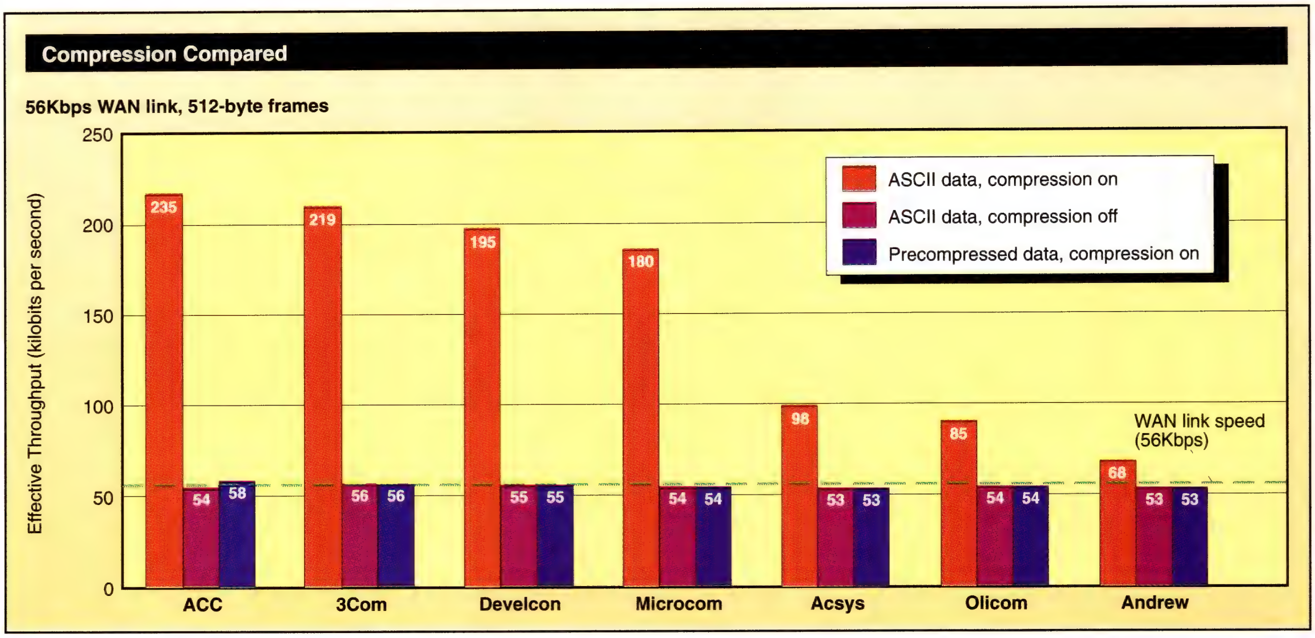

Andrew

Andrew's Pathwise 7622 has plenty of good features — ease of configuration, support for SNMP and IBM network management, and front-panel status display. Unfortunately, those pluses are all negated by the poorest data compression of any product we tested. Andrew's software-based data compression routine improved throughput by only 21%, compared with performance improvements of at least 200% posted by other products (and ACC's 400% boost).

Andrew's DOS-based configuration program, TR Bridge, was relatively easy to use. Commendably, the program offers a number of filtering options, such as include/exclude options for MAC filters and support for up to 12 NetBIOS filters. But only three SAP filters can be set, and only one group filter.

As with bridges from Acsys and Olicom, the Andrew bridge cannot be configured form a PC running NetWare; net managers must first unload the PC's IPX stack.

Develcon Electronics

Barely the size of a stackable hub, the Develcon Token Ring Remote Bridge Model 200 nonetheless combined solid performance with a rich set of features. And small size is more than a cosmetic issue, particularly at large data centres where 50 or more bridges may be deployed. Unfortunately, however, the smaller form factor means the Develcon can't be mounted in a standard 19-inch rack.

Configuration through an ASCII terminal was extremely simple; indeed, the bridge earned ratings of 9 or 10 in all of our configuration tests.

But setting some types of filters called for a little more determination. Develcon's bridge supports all kinds of filters except for Link Limiting, and users can design custom filters based on bit or byte patterns and then link the filters using Boolean operators. Doing so, however, is not a task for novices.

Develcon's technical support service, which was uniformly excellent, offered clear explanations of all bridge features. The vendor also says its next software release, now in beta testing, will add support for wildcards in filters.

Develcon was the one vendor in the test to support IBM's LAN Network Manager but not SNMP.

The Model 200 implements the same data compression technology used in Acsys's Tokenmaster 5000 — with far more favourable results. Develcon's 3.48:1 compression ratio was nearly twice that of Acsys, but a shortcoming is that the top speed for compression is 256Kbps.

Microcom

Low cost, ease of operation, good performance — Microcom's MBR/6003 offers a lot to the entry-level bridge buyer. It makes sense that Microcom's PC-based bridge would do well in our compression tests; after all, the vendor developed the MNP (Microcom Networking Protocol) compression technique that's now a standard feature on most async modems. On the downside, the Microcom bridge lacks support for key IBM features. What's more, the product is available only as cards and software loaded into a PC. The large size spells trouble for net managers having to deal with 50 or more of these units.

The MBR/6003 is actually a bridge/router, one of three such products evaluated. Fortunately, the routing features don't clutter up the bridging configuration routine, which was generally simple and straightforward. There were two exceptions, though. First, setting MAC filters means porting through several menus. Second, making any changes forced the bridge to reboot itself — with no warning. The unplanned downtime may come as an unpleasant surprise for many network managers.

In terms of features and performance, Microcom's entry is a solid one. The bridge compressed ASCII data at a ratio of better than 3:1. What's more, the MBR/6003 supports compression at up to 1Mbps — the second-fastest rate allowed. The bridge also supports dual wide area network loads between the ports.

The Microcom bridge does suffer from two notable omissions. First, it does not support IBM's LAN Network Manager, ruling out any chance to oversee it from Big Blue's network management platforms. (The bridge does support SNMP and HP OpenView.) Second, the MBR/6003 cannot filter NetBIOS traffic — essential for keeping bandwidth-gobbling acknowledgments off the wide area.

Olicom

Like Microcom, Olicom offers a PC-based bridge at an attractive price. The company's OC-3146/Remote Bridge 16/4 also supports key IBM features missing in Microcom's product. Unfortunately, it's what Olicom doesn't deliver that counts: the bridge proved more difficult than most to configure, and its data compression performance trailed that of most other entries.

Olicom's MS DOS-based configuration program, OBRIMG, and the bridge's documentation were less than a model of clarity. Some very important settings, such as enabling and disabling compression and changing the wide area network link speed, are set through command-line switches

Which publication is read by more communications professionals?

ATUG - The Exhibition and Conference of the Australian Telecommunications Users Group is unquestionably the major communications event in Australia. Visitors to the annual event are widely considered to be the region's top communications professionals, responsible for the lion's share of its networking and telecommunications purchasing.

Upon registering at ATUG'93 visitors were requested to indicate the communications/networking publications they regularly read. Judging from the **4,690*** completed questionnaires *Australian Communications* is clearly the choice of communications professionals

" I regularly read...."

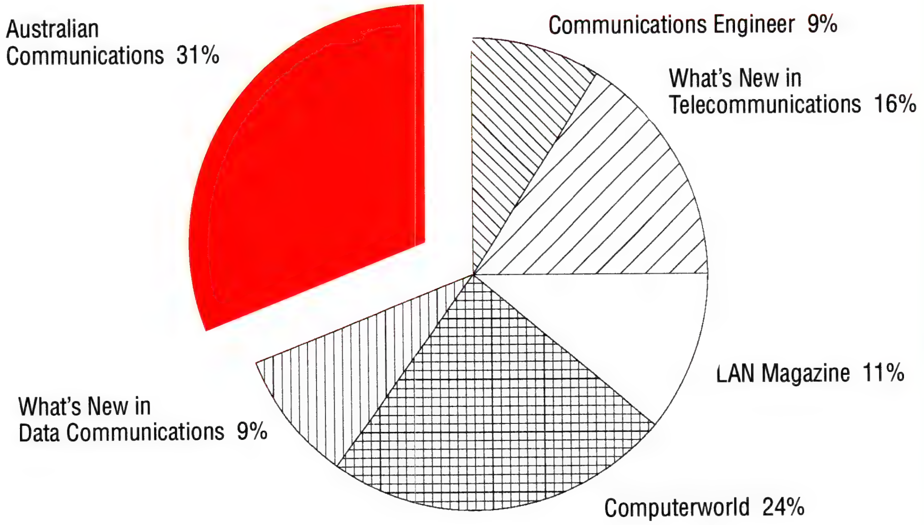

Data supplied by Riddell Exhibition Promotions, organisers of ATUG'93

* Does not include conference delegates, pre-registrations, or visitors with completed coupons from advertisements in *The Australian* and *Australian Communications.*

AUSTRALIAN
communications
The Data and Telecommunications Management Magazine

International Advertising Representatives

Head Office
AUSTRALIA & NEW ZEALAND
Australian Communications
4/541 Kent Street
Sydney NSW 2000
Australia
Ph (02) 264 2200 Fx (02) 264 2244

USA & CANADA
M&T International Marketing
Miller Freeman Inc
411 Borel Drive Avenue
San Mateo CA 94402
Ph (415) 358 9500 Fx (415) 358 9739

EUROPE
A-Z International Sales
70 Chalk Farm Road
London NW1 8AN
United Kingdom
Ph 71 284 3171 Fx 71 284 3174

JAPAN
Media Sales Japan Inc
Tamuracho Building
3-3-14 Shimbashi
Minato-ku Tokyo
Japan
Ph 3 3504 1925 Fx 3 3595 1709

TAIWAN
Acer TWP Corp
4F 19-1 Lane 231
Fu Hsing N. Road
Taipei 104
Taiwan
Ph 2 713 6959 Fx 2 715 1950

KOREA
Jes Media International
5th Floor Banpo Building
40-7 Chamwon-Dong
Seocho-Gu Seoul 137-030
Korea
Ph 2 545 8001/2 Fx 2 549 8861

HONG KONG
Acer TWP Corp HK
Unit 2, 5th Floor
82-84 To Kwa Wan Road
Kowloon
Hong Kong
Ph 764 3857 Fx 764 3857

SINGAPORE
Eastern Publishing Associates
1123 Serangoon Road #03-01
Singapore 1232
Ph 296 6166 Fx 298 7551

ISRAEL
Nora Lifton
Kfar Saba 44444
Israel
Ph 52 970065 Fx 52 953194

Test Methodology

Eleven vendors of remote Token Ring bridges were invited to participate in this evaluation: Acsys, Advanced Computer Communications (ACC), Andrew Corporation, Crosscomm, Develcon Electronics, Microcom, Newport Systems Solutions, Olicom, RAD Network Devices, Teleglobe Communications and 3Com. Four vendors declined: Crosscomm, Newport Systems, RND, and Teleglobe.

The test bed used in this evaluation reflected the most common configuration for remote Token Ring bridges: two 16Mbps Token Rings were connected over a 56Kbps wide area network link by a pair of bridges (see the figure).

We used a null modem from VIT Inc. to simulate the WAN link. A Novell NetWare server was attached to one Token Ring, while a NetWare client was attached to the other ring.

To measure bridge throughput, traffic was generated on one LAN and sent across the WAN link to the other LAN, where data rates were measured with a protocol analyser. Three kinds of traffic were used: first, we established a performance baseline by sending ASCII data with bridge compression disabled. Second, we sent the same ASCII data with compression turned on, and noted the difference in effective throughput. Finally, we sent precompressed data through the bridge to verify that the bridge's compression algorithms would not degrade throughput by attempting to compress uncompressible data.

To create the data, the NetWare client copied a 1-megabyte file to the server; this exchange was captured with an Expert Sniffer from Network General and then edited for use on the Framethrower frame generator from LANquest Labs. We used a stream of 20 unique frames for each of our tests, each with a length of 512 bytes, including header. Compressed data was created using PKzip, a data compression program from Pkware. Throughput was measured with an Expert Sniffer and with a DA-30 protocol analyser from Wandel & Goltermann.

We used the Tolly Group's tried-and-true 'blip' test to gauge bridge latency. In this test, a DA-30 is connected to both the sending and receiving LANs. The difference between the timestamps is considered the latency.

Kevin Tolly and David Newman

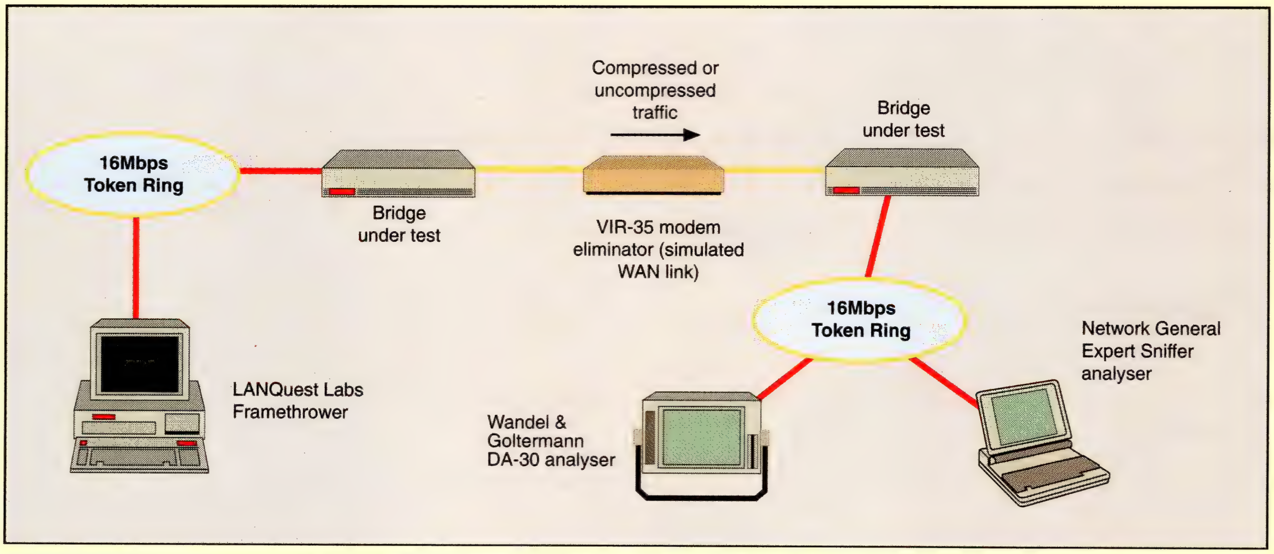

when launching OBRIMG rather than through the configuration program itself. And the Olicom bridge, like products from Acsys and Andrew, cannot be configured from a personal computer running Novell NetWare.

We were also fairly underwhelmed by the bridge's performance. Olicom's bridge boosted effective throughput by about 50% — not much of an improvement considering that first-place ACC posted a gain of more than 400%.

There is some good news out there for Olicom customers. The vendor is now shipping its Token Ring Remote Bridge OC-3445, the first Token Ring bridge, local or remote, to offer an intelligent CAU (controlled access unit) module. This two-in-one deal eliminates the need for separate units. Olicom also plans to offer a router module for the 3445. Initial 3445 tests show data compression to be no higher than that of the OC-3146/Remote Bridge 16/4, but users may be intrigued enough by the dual design to forgo the advantages of higher compression.

3Com

3Com's Netbuilder Token Ring has basically the same characteristics as most of the vendor's numerous other products: excellent performance and configuration routines you either love or hate.

The 3Com Netbuilder TR is one of three bridge/routers featured in this test. On the plus side, few devices support more protocols, transported over more kinds of links, and filtered more kinds of ways, than the Netbuilder TR. But all those choices may prove overwhelming, especially to network managers who want to deploy the Netbuilder TR solely as a remote Token Ring bridge.

But provided network managers find their way through Netbuilder's myriad and extensive configuration routines, they'll be rewarded with first- class performance. The Netbuilder TR compressed ASCII data by a ratio of nearly 4:1 — the second-best-showing in the test. One drawback, however, is that compression is only available on 56 or 64Kbps leased lines, even though the product supports wide area network links as fast as 9.22Mbps.

Like several other bridges in this test, the Netbuilder TR can't be managed by IBM's LAN Network Manager platform, a problem for users in all-IBM environments. 3Com has pledged LNM support by the fourth quarter of this year.

Kevin Tolly is Director of the Data Comm Test Lab and President of The Tolly Group. He is based in Sea Girt, New Jersey. David Newman is Testing Editor for Data Communications magazine, based in Boston. InterLab research associate Karl Metzelaar handled all testing for this story.

The next time your network needs a more advanced hub, you'll already have it.

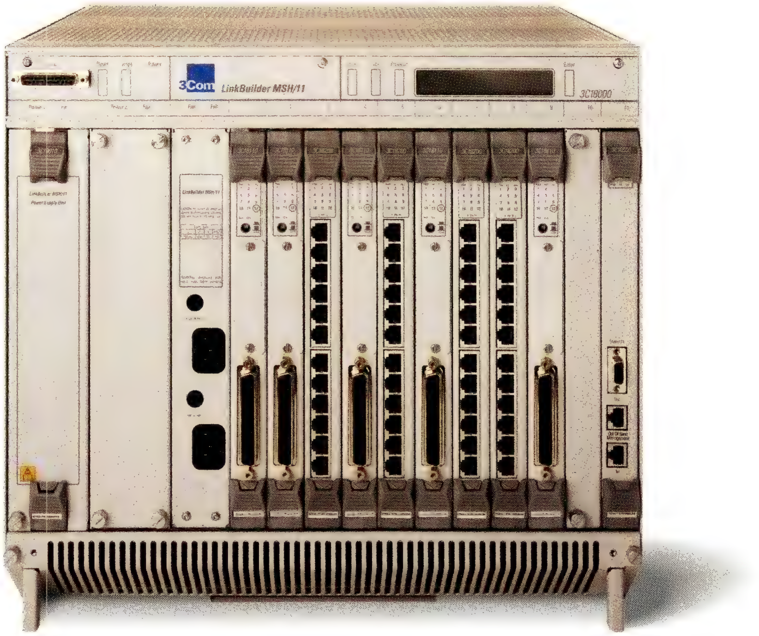

For a truly intelligent way to network, why not choose a hub advanced enough to be its own replacement? A hub designed to easily accommodate any change your growing network requires. A hub flexible enough to allow you to customize and configure your network as it grows.

It's 3Com's LinkBuilder® Multi Services Hub. The most advanced, cost-effective networking solution in the industry.

Thanks to its unique "technology-independent" backplane, rather than adding backplanes for every new technology, the LinkBuilder MSH™ allows you to mix Ethernet and Token Ring networks today, and add ATM and other new networking technologies when you need them. All the while avoiding costly and disruptive hardware upgrades.

But why not see for yourself. Just call us a call at **(02) 959-3020, ext. 20** for a free copy of 3Com's white paper "High-Performance Scalable Networking," and more information about the LinkBuilder MSH. You'll learn the easiest way to plan for the future is to have a hub that's already there.

3Com®

Networks That Go the Distance™

©1993 3Com Corp. (408) 764-5000. 3Com, Networks That Go the Distance, MSH, and LinkBuilder are trademarks of 3Com Corp.

Anixter Australia	(02) 333 0800	Bull HN Information Systems	(02) 847 7100	Fujitsu Australia	(02) 410 4555
GEC IT	(02) 415 0500	Interconnect NZ	(649) 358 0322	LANCON	(02) 899 3033
NetArch	(07) 393 1933	NETSOURCE	(09) 470 5311	NST	(02) 499 2900
Scitec	(02) 428 9555	Tech Pacific Australia	(02) 697 8666	Tech Pacific NZ	(649) 415 8300

Are dial-up problems with 3270, AS/400 or LANs letting you down?

For more than 10 years now, Rosser Communications has been working to solve communication and connectivity problems for business. We provide a comprehensive range of data comms solutions and products.

No matter what host or workstation configuration you have, we can provide the experience and equipment so that you can connect with confidence through dial, leased or X.25 networks.

Rosser Communications, your effective data comms solution.

Call Alex Rosser or Kendel Bott on (02) 418 2544 and solve your data comms problems today!

ROSSER
COMMUNICATIONS

Rosser Communications Pty Ltd ACN 003 021 841
Suite 3, 6 West Street, Pymble NSW 2073
Ph: (02) 418 2544 Fax: (02) 418 2501

Multimedia Switching Platform

AWA has announced what the company claims is the first truly integrated digital multimedia switching platform providing distribution and integration of digitised voice, video, synchronous and packetised LAN data.

The system, known as AWAnet, is implemented as centralised or distributed switching hubs. Platforms now available include 6- and 20-slot modular chassis and PC-based systems.

AWA officials said specific video, voice or data switching requirements are fulfilled using a mix of hub controller cards, network interface cards and application software and configuration tables downloaded from an SNMP-based network management station.

The system's modular chassis provides a multi-gigabit-capable backplane architecture consisting of multiple Ethernet, Token Ring, isochronous N x 64 (up to 400Mbps) and cell switched buses. Redundant power supplies are included, and all modules are hot swappable. WAN access is provided by ISDN primary and basic rate, G.703, E1/ T1, analogue exchange lines and frame relay. Distributed PABX switching is supported between chassis remotely connected via G.703/704, or via FDDI II in the LAN campus.

Models available include 12-port 10Base-T Ethernet, 6-port 10Base-F, 5-port multiprotocol router and bridge, 12-port Token Ring lobe, 4-port fibre Token Ring RI/RO, 4-port telephone exchange, 2-port E1/ T1, 4-port sync and FDDI II access controller. Officials said models currently under development include ATM SDH with AAL1 to AAL5, switched Ethernet, fast Ethernet, isochronous Ethernet, NOS/application server, voice compression, ADSL and wireless interfaces.

Management of AWAnet systems is via an SNMP console, which is available for both DOS/ Windows and Unix platforms. Adaptor systems are available for third party SNMP consoles including Hewlett-Packard OpenView, SunNet Manager and IBM NetView/6000.

AWA (02) 887 7374

3Com Fine-Tunes Remote Access

3Com Corporation has released new internetworking software and hardware which the company says will aid WAN access to remote sites, increase data prioritisation options and automate centralised network management operations. The announcements follow 3Com's recent introduction of its NETBuilder Remote Office product line.

The company says users will now have true interoperability with the Novell Multiprotocol Router Plus 2.X and other routers using IPXWAN (RFC 1362) internetwork Packet Exchange for Novell WAN environments, as well as non-meshed frame relay and dial-on-demand support to peg rising WAN costs.

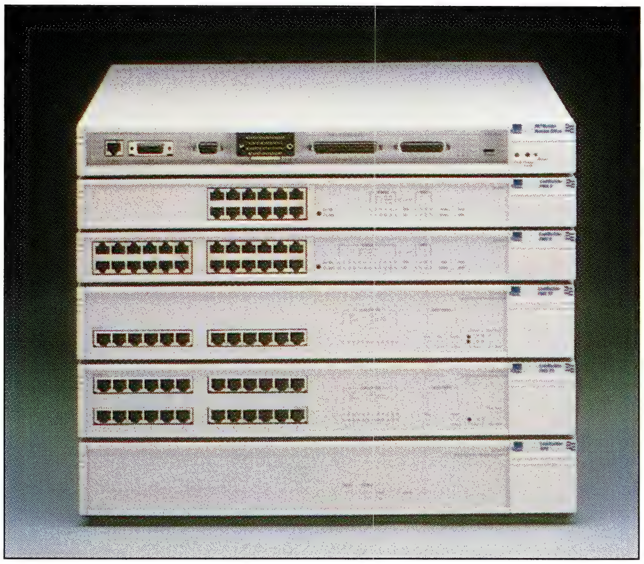
3Com's new internetworking equipment helps stem rising WAN costs

Officials said that users will also achieve capital savings on costs associated with WAN operations via a new high-speed V.35 tri-port module which effectively triples WAN port density.

IPXWAN support over Point-to-Point Protocol (PPP) can provide users with router-to-router negotiation before RIP and SAP updating, said officials, adding that network managers will have greater flexibility on how IPX SAP and RIP updates are broadcast over WAN links.

A four-tier data prioritisation structure and new queuing arbitration algorithm ensures no queue is deprived of WAN access, and that urgent/high priority traffic is sent with the minimum delay.

3Com (02) 959 3020

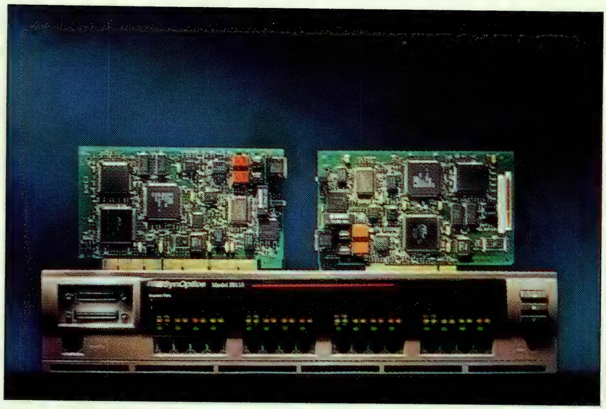
The LattisSwitch range can handle up to 18 100Mbps connections

SynOptics Fast Ethernet Switching Hub

SynOptics has announced the release of a high performance family of Fast Ethernet switching hubs which the company says will boost users' existing network bandwidth by up to 200 times.

The new hub family, known as the LattisSwitch System 28000, is based on SynOptics' 2Gbps Fast Frame technology and the 100Base-X specification developed by the Fast Ethernet Alliance. The 28000 family offers both 10Mbps and 100Mbps Ethernet connections, and, unlike existing shared media LANs, allows each attached dedicated or multiple station to transmit simultaneously, resulting in higher overall available bandwidth.

The hubs are able to handle up to 18 switched 100Mbps connections, and officials said the new products also offer a lower-priced alternative to routers when used to break up large networks into multiple segments to improve performance.

They work with all existing 10Base-T cards, and are also fully interoperable with Intel's newly-released 10/100Mbps EtherExpress PRO/100 LAN adaptor cards, which are the result of a joint technology relationship between the two companies.

The LattisSwitch range will be available in four versions, which will be rolled out over the next six months. The first, the Model 28115, is available from this month priced at $31,475.

SynOptics has also announced new switched 10/100Mbps virtual LAN management software for Ethernet networks in version 5.0 of its Optivity network management system. The virtual LAN capability gives allows users to dynamically configure mixed shared-media and switched networks using an object-oriented user interface.

SynOptics (03) 853 0799

New MMAC Module

Cabletron Systems has introduced the new TPXMIM, a port assignment module for its Multi Media Access Centre (MMAC) family of intelligent hubs. The TPXMIM uses the Cabletron

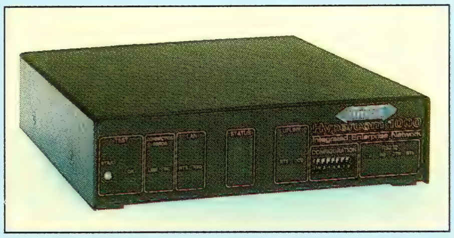

The IEN 1000 integrates a range of branch communications functions in a single box

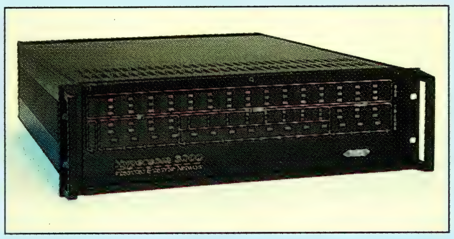

The Hypercom 6000 manages the connection of a large number of branch offices to a central site

New Australian Designed Integrated Branch Communications Solutions

Australian manufacturer Hypercom has released two new products designed to serve the needs of companies that need to connect branch office sites to the corporate network.

The first is the Hypercom IEN 1000, which the company says is the industry's first 'branch-in-a-box' communications system.

The IEN 1000 integrates intelligent hub, router and all key branch office communications functions, simplifying installation, maintenance and management while at the same time improving reliability and lowering costs, according to company officials. It combines Ethernet and Token Ring intelligent hub and router functions, legacy protocol support, dial-up backup modem and data encryption and compression. It incorporates Hypercom's unique SNA/LAN consolidation technology which combines SNA and LAN traffic without the need for tunnelling or encapsulation, into a single data stream that can be transported over public or private networks. Officials said this technique both protects existing investments in SNA equipment, and greatly reduces leased line costs by eliminating the need for separate parallel networks for handling SNA and LAN traffic.

The integration of several devices in the one unit overcomes interconnection failures between different boxes, which are responsible for a large percentage of communications failures on networks. This is of even greater relevance in a branch office environment, which may not have technical staff on hand to solve network problems.

The unit can support up to 24 LAN workstations at a branch and can combine multiple synchronous protocols, bridge/router and hub functionality and non-routable LAN protocols as well as asynchronous communications, X.25 and other protocols. The WAN interface supports frame relay and public and private networks. The IEN 100 can also incorporate high-speed modems, E1 interfaces, RS-232 and V.35. Other features include integrated protocol conversion (SDLC to LLC-2), data compression (Lempel-Ziv), DES data encryption and fully automated dial backup.

The IEN 1000 base model is configured with one WAN port supporting data rates of up to 1.5Mbps, including frame relay; one RS-232 port supporting SNA, SDLC, async, bisync, X.25 and other legacy protocols. Options include fully managed Ethernet and Token Ring router and 4-port hub modules, expansion to four RS-232 ports, V.32 dial backup modem, data compression, DES data encryption and a redundant power supply for extra reliability. A typical configuration with one WAN interface, an Ethernet router, one serial interface for SNA and one 4-port Ethernet hub is priced at $5,500.

Additionally, the company has released two multiprotocol backbone nodes, the Hypercom 6000 and 7000. The nodes allow a large number of IEN 3000 branch nodes to be concentrated at host sites, allowing organisations to cost effectively and efficiently link hundreds of branch terminals, PCs and telephones to the corporate headquarters' mainframe, servers and PABX.

Like the IEN 1000, the Hypercom 6000 and 7000 use the company's own system to consolidate LAN protocols and SNA, and offer many of the same protocol handling features. The products support RIP and OSPF, and can use all possible routes within the network and build their own routing tables.

Pricing for a typical configuration with 8 WAN interfaces, 4 dual LAN interfaces and 16 serial interfaces starts at around $30,000.
Hypercom (02) 956 6066

multichannel hub architecture to provide port assignment capabilities in an integrated bridging and routing environment.

The module requires only a single management module to monitor and control up to 168 ports with full RMON services, and allows users to be grouped based on file sharing activities or network usage, thus maximising bandwidth utilisation.

Via port assignment, users on a single hub module can be distributed to any of the three Ethernet networks on the MMAC backplane, and the TPXMIM also allows for board assignment, permitting the assignment of an entire module to a single Ethernet channel.

When used in conjunction with EMM-E6, users have the option of creating an uplink to ATM, FDDI, Token Ring, Ethernet or a WAN connection, via Cabletron's BRIM bridge/router interface module.

TPXMIM is manageable using Cabletron's Spectrum for Open Systems for both Unix and DOS, or any SNMP-based network management system. It also includes Cabletron's Lanview LED diagnostic indicators to provide status information about the module, including Link, Receive Data, Collisions and Network Configuration.
Cabletron (02) 950 5900

New DEC Servers

Digital Equipment Corporation has announced new powerful severs and software which officials say will aid the transition to client/server environments.

The new products include the Digital 2100 Server Model A500MP symmetric multiprocessing (SMP) Alpha AXP server, the AdvantageCluster program, TechAdvantage parallel application development tools, enhanced operating system software and a range of new technical applications.

The Digital 2100 Server Model A500MP is the first in a family of RISC-based departmental servers combining SMP with low cost PC technology. The servers will support over 5,000 currently available Alpha AXP applications, said officials.

The system is based on the 190MHz DECchip 21064 processor. It is expandable from one to four processors, with pricing from $37,800 for a one-processor system to $67,200 for a four-processor configuration. Users can choose from DEC OSF/1, Open VMS AXP or Microsoft Windows NT Advanced Server software for the system.

Digital's AdvantageCluster for Unix initially will comprise three packages:

- AdvantageCluster Compute Servers, cabinet-mounted arrays of four to 32 DEC 3000 AXP workstations with Ethernet or high-speed GIGAswitch interconnect;
- AdvantageCluster File Server, which delivers very high NFS throughput (up to 4,000 NFS operations per second); and
- AdvantageCluster Available Server, which uses clustered AXP systems in a high-availability configuration.

The hardware announcements are accompanied by a new version of DEC's OSF/1 operating system with SMP support and other performance-boosting features such as dynamic device driver loading, standard-compliant real-time system services and a bundled DEC Ada Run-time Library

A new version of DEC's Open VMS operating system is also to be released, and the company additionally has announced a range of new and updated systems and network software for workgroups and departments including the VAX 4000 Model 105A and the MicroVAX 3100 Model 95, InfoBroker software, PathWorks Software Development Kit; DECMessageQ version 3.0; Digital Distributed Computing Environment; and DEC/EDI version 2.0.

DEC (02) 561 5252

New Wireless Family

RangeLAN2 offers the longest range of any product in its class

PowerCorp has announced the release of the Proxim RangeLAN2 family of Spread Spectrum wireless LAN adaptors. The new range includes an ISA model, a PCMCIA model and a wireless local bridge.

The RangeLAN2 series is based on frequency hopping technology in the 2.4-2.283GHz band. Officials said the product offers the longest range of any wireless product in its class, operating at a distance of up to 300-500 feet in normal office locations and up to 800-1,000 feet in open space environments.

The RangeLAN2 PCMCIA adaptor, scheduled for release later this year, enables portable computer users to access existing wired client/server networks or to communicate on a peer-to-peer basis with other RangeLAN2 systems. The product features low power consumption (100 milliwatts), and supports PCMCIA 2.0 software with 'hot insertion' capability.

The RangeLAN2 ISA adaptor fits into the expansion slot of most desktop PCs or servers, and is well suited to providing a wireless extension to conventional wired LANs in places that are too difficult or too costly to wire. It can also provide a standalone peer-to-peer wireless solution for small workgroups.

The RangeLAN2 wireless bridge enables transparent access to a standard wired LAN. It operates at the MAC layer of the OSI network model, and allows RangeLAN2-equipped portable and desktop computer users to achieve operating system and protocol-independent wireless access to the wired LAN. The product currently supports Ethernet LANs only, but a Token Ring version is planned.

PowerCorp (02) 476 3466

3Com AccessBuilder

3Com has announced internetworking and security enhancements to its AccessBuilder family of remote access servers.

New features include dynamic IP address assignment and better OSF/DCE security. Regardless of their location, users requiring multiple site access to the corporate LAN can take advantage of a 'floating IP address pool' which makes use of a limited number block of dedicated IP addresses.

AccessBuilder uses a floating IP address pool for remote access

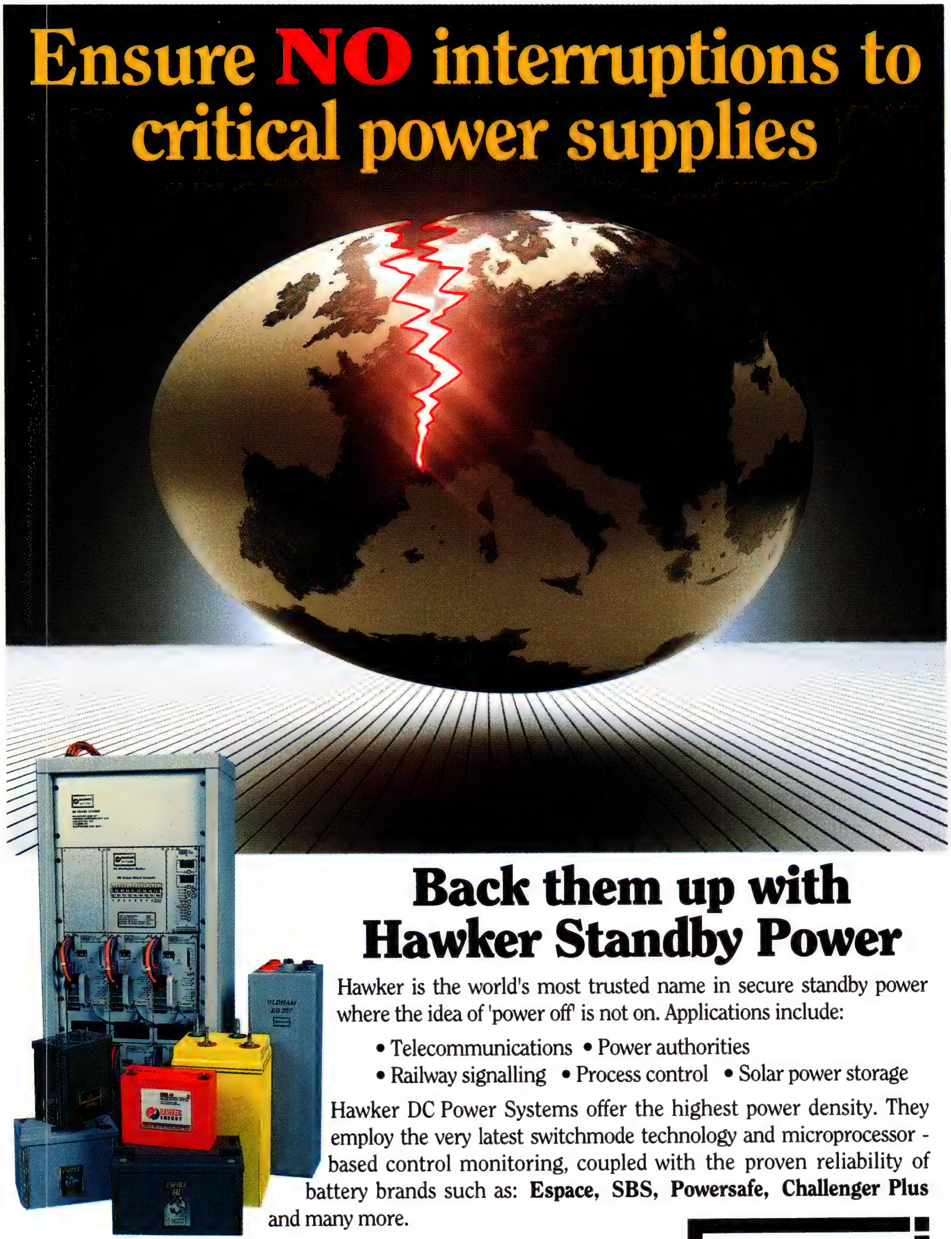

Ensure NO interruptions to critical power supplies

Back them up with Hawker Standby Power

Hawker is the world's most trusted name in secure standby power where the idea of 'power off' is not on. Applications include:

- Telecommunications
- Power authorities
- Railway signalling
- Process control
- Solar power storage

Hawker DC Power Systems offer the highest power density. They employ the very latest switchmode technology and microprocessor - based control monitoring, coupled with the proven reliability of battery brands such as: **Espace, SBS, Powersafe, Challenger Plus** and many more.

For details of the superior response, reliability and cost-efficiency of low-maintenance Hawker Power Systems, please phone or fax for comprehensive technical information.

SYDNEY: 119-127 Wicks Rd North Ryde NSW 2113 Tel: (02) 878 5000 Fax: (02) 878 5555
MELBOURNE: 832 High St East Kew VIC 3102 Tel: (03) 810 9568 Fax: (03) 810 9589
QUEENSLAND: 240 Ipswich Rd Buranda QLD 4102 Tel: (07) 391 0444 Fax: (07) 391 0400

HAWKER BATTERIES

POWER BY DESIGN

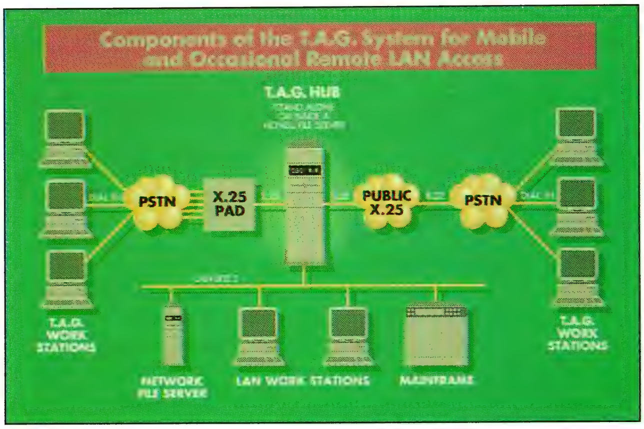

TAG connects remote users to corporate LANs over the PSTN

OSF/DE security allows network administrators to centrally manage security via a security server, which acts as an agent between the AccessBuilder and the network security server.

3Com officials said newly automated client-side software maximises WAN access efficiency and adds a layer of security and ease of administration through features such as Dial Scripting and Chaining. Dial Scripting enables users to support third party front-end security devices and to store scripts for easier access to commercial services such as CompuServe. Chaining allows remote access users to pre-establish the configuration of their working environment before dialling into the corporate network.

The AccessBuilder products are based on 3Com's Personal Routing system architecture, the personal office internetworking equivalent of 3Com's Boundary Routing system. Through network node emulation, the Personal Routing system provides individual remote users with full multiprotocol access to resources on the enterprise.

The AccessBuilder family supports data reduction and compression, and each of the 16 ports can be individually configured to support a range of WAN services. Pricing starts at $6,025 for 4-port models, $7,745 for 8-port models, and $10,849 for 16-port models. Pricing for AccessBuilder Token Ring products starts at $7,745 for 4-port models, $9,470 for 8-port models, and $12,570 for 16-port models. **3Com (02) 959 3020**

New Remote Access System

Melbourne-based firm Network Designers has recently launched the Transparent Access Gateway (TAG) product, which company officials say offers a low cost way for remote workers to access corporate networks over regular telephone lines.

To use the system, users only require TAG and a Hayes-compatible modem, making the system especially suitable for workers who only require occasional access to the LAN, according to company officials.

The TAG system has two major components. The TAG Hub runs on the LAN as either a standalone application or inside a Novell file server. It provides X.25 links to external PCs using an X.25 Packet Assembler Disassembler (PAD) or Telecom's Austpac public X.25 service. The TAG Hub supports 64 concurrent remote sessions.

TAG allows a remote PC user to appear as a LAN node on a remote network via an asynchronous dial-up line. Dial-up access is automated via the software, avoiding the need for add-in cards and making it especially suited to laptop and notebook PCs.

The TAG workstation runs under DOS or Windows, and

Open Wiring Systems Pty Limited ACN 003 658 857

N.S.W.
28 Punch St, Artarmon 2064
Ph (02) 906 5777 Fax (02) 906 2777

A.C.T.
26 Dunbar St, Fraser 2615
Ph (06) 258 0090 Fax (06) 258 0081

QLD
Suite 6, 30 Argyle St, Breakfast Creek 4010
Ph (07) 862 4100 Fax (07) 862 4144

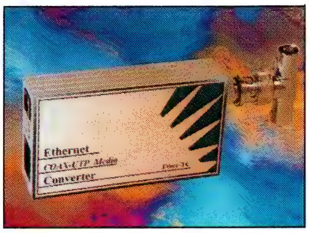

The Coax-UTP Converter is only slightly bigger than a matchbox

depending on hub configuration, can communicate with an Ethernet LAN at speeds of up to 28.8Kbps, or even higher with compression. It utilises Network Device Interface Specifications (NDIS) which allow application programs to deal with remote workstations as if they were fitted with a standard Ethernet card supporting the NDIS interface. The interface also supports other protocols such as TCP/IP and IPX.

The software can automatically initiate calls dependent on the destination Ethernet address, TCP/IP network address or NetBIOS name matching, using call-making scripts written by the network administrator.

Network Designers KNX (03) 629 3800

Coax-UTP Converter

Speedycom has released its new Coax-UTP Converter, which provides one UTP port for connection to a twisted pair cable segment, and one BNC port for connection to a thin coaxial cable segment.

The self-powered module is easy to install, and suitable for mixed media environments, according to company officials, who added that it is especially useful for connecting a small thin Ethernet workgroup to the main 10Base-T network.

The new converter can save the cost of installing both an extra 10Base-T hub and multiple segment cabling by connecting groups of stations which are located far away from a concentrator. It can also be used for Ethernet connections to twisted pair cabling where BNC and UTP ports are not provided.

The unit measures 8cm by 4cm by 2cm (slightly larger than a matchbox), and has LEDs for reception, collision, jabber, power and link. It is guaranteed for two years, and is priced at $265, excluding tax.

Speedycom (09) 388 1755

Spectrum 3.0

Cabletron has unveiled version 3.0 of its distributed enterprise network management platform, Spectrum. The new release extends the scope of the package to include wide area networks, telephony and PABX management as well as ATM and connection oriented systems.

Officials said that by distributing Spectrum's advanced modelling engine over multiple processors in a client/server architecture, network management traffic is reduced while management capabilities such as security and redundancy are significantly enhanced.

The most feature-rich of the company's range of network management applications, Spectrum 3.0 is based on distributed databases called SpectroServers. Spectrum uses inductive modelling (a form of artificial intelligence), allowing management capabilities to be distributed over multiple Spectro-Servers while one or more Spectro-Graph graphical user interfaces, acting as a client, allows for navigation through the network.

New management functions include: Spectrum Configuration Manager, which lets users change and monitor the configuration of SNMP-based system and network elements; Report Generator, which supports six different report types; Auto-Discovery, which automatically locates the device models and places them in view in order to update the distributed network management databases; and a point-and-click Control Panel; Enhanced Fault Isolation Performance, which can decrease the time need to isolate network faults by up to 400%.

Spectrum 3.0 runs under all major versions of Unix and will be ported to Windows NT later this year.

Cabletron (02) 950 5900

Branch Office Solution

Novell has announced the release of its new NetWare Multi-

LM1 PocketScope

The world's most portable PC-based WAN Analyser.

Special Offer
Save 20% on orders
before 30th June

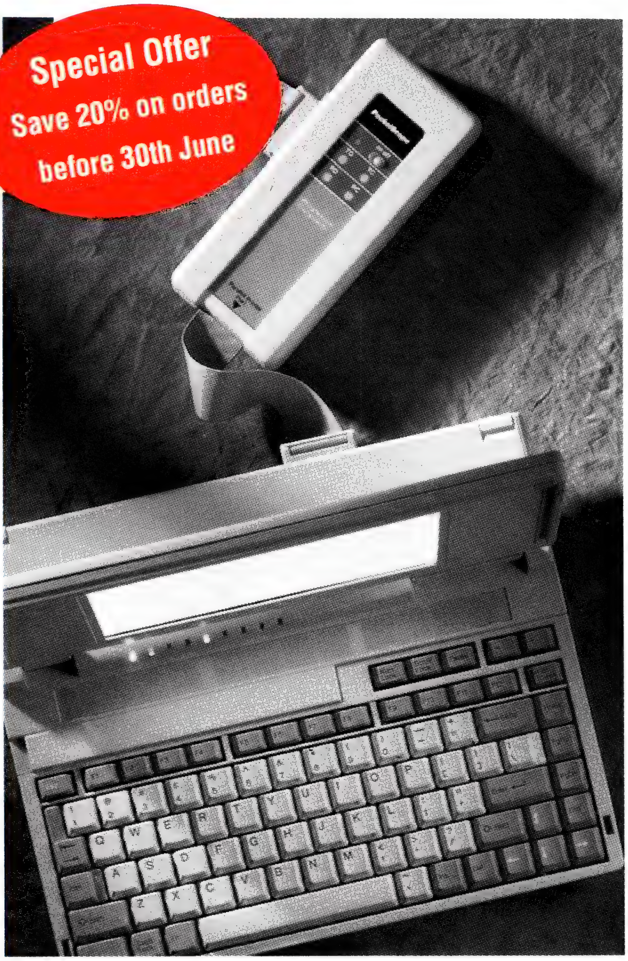

A fully featured protocol analyser that connects to the **parallel port** of any PC...(No expansion slots required).

Plus

- SNA, X.25, Q.931, QLLC, DDCMP, full decodes and Level 2 real time statistics.
- Pre- and Post capture filtering.
- Extensive library of automated test routines for SNA, X.25, ASYNC and BISYNC.
- Programming language.

TESTCOM-DATA....MORE CHOICE

TESTCOM-DATA
217 North Rocks Road, North Rocks NSW 2151
Telephone: (02) 630 7528 Fax: (02) 630 7226

protocol Router 2.11 and Multiprotocol Router Plus 2.11 software, which has been designed for local network routing and wide area communications in branch office environments.

Frame relay support has been added to provide a cost effective solution for the routing of IP, IPX, AppleTalk and source-route bridging traffic.

Additionally, data compression capabilities can more than double throughput over leased line connections, according to officials, and a remote installation facility has been added.

The MPR 2.11 and MPR Plus 2.11 also have new filtering capabilities to help safeguard data broadcasts and isolate certain branch traffic, so users can selectively control the amount of data traffic sent across links. The products are highly integrated with NetWare 3.12, and have a number of other new features including: IPX Routing Information Protocol/Service Advertising Protocol (RIP/SAP) filtering; PPP compression; and the ability to route and source-bridge multiple protocols including IPX, IP, AppleTalk and OSI locally, and remotely over PPP links, frame relay and X.25 networks.

The products are priced at $1,835 for first-time buyers, and there are upgrade deals for existing users.

Novell (02) 925 3000

Pocket Fax Modem

The Ibislink is both a 14,400bps modem and a G3 fax

Alloy has introduced the new Ibislink pocket-sized modem with fax functions. The Ibislink operates as both a 14,400bps modem and a 9,600bps fax with V.42*bis* 57,600bps throughput.

For modem applications, the Ibislink is compatible with CCITT V.32*bis*, V.32, V.22*bis*, V.23, V.22, V.21, Bell 212A and Bell 103 standards. The unit also incorporates features such as flow control and speed conversion including MNP2-4 and V.42 Error Correction.

Ibislink uses both MNP5 and V.42*bis* compression schemes while operating in asynchronous mode, allowing it to achieve an effective throughput of up to 57,600bps with automatic fallback capabilities. Because Ibislink users the extended Hayes AT command set, it is compatible with most other communications packages, said officials.

As a fax, Ibislink functions as a G3 system compatible with CCITT V.29 and V.27*ter*, and has the ability to send and receive faxes at 9,600bps using class 2 fax command sets.

The unit comes with its own power adaptor, but can also run off a 9V battery to increase portability. The entire unit weighs around 180g including battery, and also comes with its own carry case, cables, and software for fax and modem operation.

It can store up to 10 telephone numbers in non-volatile memory, is Austel approved, and is covered by a 12-month warranty.

Alloy (03) 574 9891

Enterprise Backup

Memorex Telex has announced a strategy called Enterprise Network Storage designed to deliver simple solutions for data backup and restore across an enterprise.

The company's first offering in the new strategy is a complete hardware, software and service solution for full enterprise backup of local area networks. The system comprises Memorex Telex's 9400 Enterprise Communications family of channel-attached and remote hardware platforms built on Novell's Net-

Overvoltages can destroy sensitive electronic systems

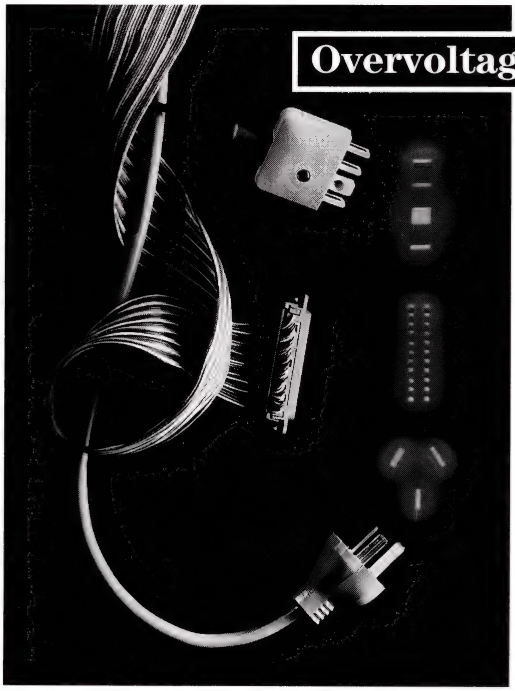

Recent advances in microprocessor technology have increased the power of the silicon chip, reduced its physical size and increased its susceptibility to damage or destruction from power overvoltages, spikes or complete power failures.

The consequences of system failure or data loss means the inevitable cost in down time and loss in productivity.

To provide reliable power to your electronic equipment, Critec offers you a choice from an extensive range of on-line and off-line UPS. From 500VA to 3kVA all units have increased surge ratings and are specifically tailored for Australia's adverse lightning conditions.

Critec's engineering experience has been recognised internationally and the Dataguard DG12 UPS has won an Australian Design Award for engineering excellence.

To find out how to protect your operational reliability, simply fill in the form below and either mail it to Critec head office or fax it back to us on **002 73 0399**.

CRITEC
Fax 002•730399 Phone 002•730066
Critec Pty Ltd Box 536 GPO Hobart 7001 TAS
NSW (02) 688 4528 VIC (03) 696 1633 ACT (06) 257 3055
QLD (07) 849 5077 SA (08) 364 2555 NT (089) 41 3880
WA (09) 257 1065 NZ (09) 309 2464

☐ **YES.** Please send me information on the Critec UPS range of reliable backup power supplies and your publication note CRPN03 "Power conditioning – an overview".

NAME ____________________
TITLE ____________________
COMPANY ____________________
ADDRESS ____________________
CITY __________ POSTCODE ______ PHONE __________ FAX __________

CRC-B029 OZ

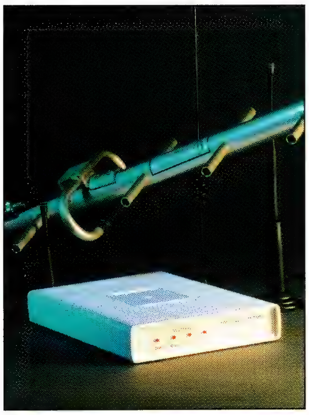

The Dataspeed provides wireless connection for up to 20km

Ware for SAA; Emprise Technologies' Stage3 Virtual Tape Solution for LANs distributed storage management software; and Cheyenne Software's ARCserve for NetWare.

The Stage3 software provides for the centralised storage and management of distributed data while at the same time allowing the LAN backup process to continue to be administered and controlled from the LAN.

According to officials, unlike mainframe-centric backup products, Stage3 is able to preserve and capitalise on existing LAN backup products while at the same time providing flexibility of configuration between LAN and mainframe resources. Stage3 can be used as a long-term archival solution to back up mission critical files, as a centralised solution which completely eliminates local media handling, or in conjunction with short-term tape backups.

Officials say Memorex Telex's 9400 Enterprise Communications family provides a wide range of connectivity options, including FDDI, Token Ring, Ethernet, SDLC and router interfaces. They said that all 9400 products provide gateway support for NetWare, SNA and TCP/IP devices.

A number of options are available to suit different environments, and Memorex Telex will provide worldwide service and support.

Memorex Telex (02) 805 5802

Radio Modem

Australian vendor Smart Radio Modems has released its new Dataspeed radio modem, which uses spread spectrum technology to provide high speed point-to-point or multi-point wireless links over distances of up to 20 kilometres.

The modem is software configurable, and once set is unaffected by power drop-outs or batteries running down. It supports asynchronous data rates of up to 115Kbaud. Company officials said that, in addition to the wireless networking of PCs, the new modem is suitable for applications such as interconnection of multiplexers or wireless links to remote data sensors.

Smart Radio Modems
(064) 525 555

NetComm Price Cuts

NetComm has announced it is reducing prices by up to 50% across its full range of Australian designed and manufactured modems. Pricing for the NetComm In FaxModem 24 and the Pocket FaxModem 24 now starts from $199, while the recently-released SmartModem M11F, which runs as 28,800bps now sells for $999.

NetComm (02) 888 5533

ISDN Bridge

AWA has announced the release of two new products in its Teleworker range. The Teleworker 200 ISDN bridge is a compression ISDN bridge for remote single users or workgroups, and can move two megabits of data traffic in two minutes across 128K of ISDN basic rate access. The Teleworker 150 is a cheaper version for single users who do not require data compression.

The Teleworker 200 features data compression, flash memory for easier upgrades, and non-volatile memory to save configuration data.

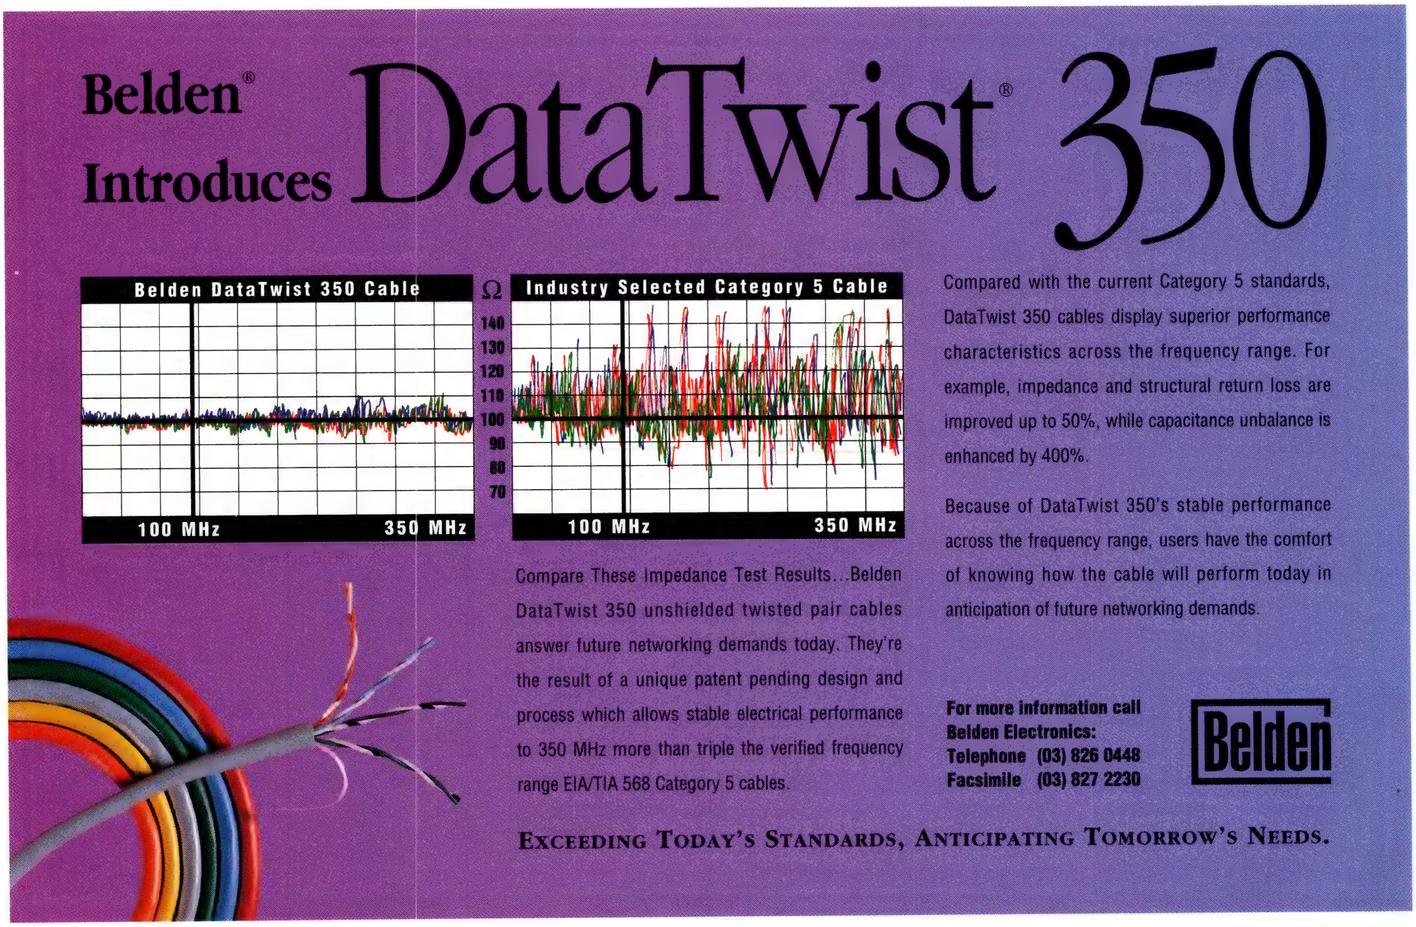
Belden DataTwist 350 Cable
Ω
140
130
120
110
100
90
80
70
Industry Selected Category 5 Cable
100 MHz
350 MHz
100 MHz
350 MHz

Belden® Introduces DataTwist® 350

Compare These Impedance Test Results...Belden DataTwist 350 unshielded twisted pair cables answer future networking demands today. They're the result of a unique patent pending design and process which allows stable electrical performance to 350 MHz more than triple the verified frequency range EIA/TIA 568 Category 5 cables.

Compared with the current Category 5 standards, DataTwist 350 cables display superior performance characteristics across the frequency range. For example, impedance and structural return loss are improved up to 50%, while capacitance unbalance is enhanced by 400%.

Because of DataTwist 350's stable performance across the frequency range, users have the comfort of knowing how the cable will perform today in anticipation of future networking demands.

For more information call
Belden Electronics:
Telephone (03) 826 0448
Facsimile (03) 827 2230

Belden

EXCEEDING TODAY'S STANDARDS, ANTICIPATING TOMORROW'S NEEDS.

The NDC PCMCIA adaptor connects portable to a LAN or host

The Teleworker 150 offers a LAN interface as a 10Base-T hub, while the Teleworker 200 features 10Base-T and 10Base-2 Ethernet ports and 10Base-T port configurable as a hub.

AWA (02) 887 7374

Token Ring PCMCIA Adaptor

IDE Technology has announced the release of NDC's new 16/4 Mbps 16-bit Token Ring PCMCIA Credit Card Adaptor. Company officials said the adaptor enables any portable computer with a PCMCIA slot to connect directly to a Token Ring LAN or an IBM host system.

The adaptor uses IBM's Token Ring mini-module to ensure IBM compatibility and interoperability. Other features include a diagnostic program with self-test function, and suitability for both Type I and Type III connectivity. The adaptor is also software configurable and complies with PCMCIA standard 2.0.

IDE Technology
(02) 899 9900

X.500 Directory

Datacraft has developed and released an X.500 Directory which lets organisations manage, manipulate and exchange information internationally regardless of host system environments or network platforms.

The product is called DX500 OpenDirectory, and Datacraft officials said it is the first X.500 Directory system in the world to use an industrial strength ANSI-standard SQL relational database as its platform.

DX500 OpenDirectory can be seamlessly integrated into current host and networking environments, and officials said it is simple to customise, has a user friendly Windows interface, and is virtually unlimited in its size and structure.

The product allows database updates to be performed at any time without affecting normal system operation. It fully supports the Directory Access Protocol (DAP) and all X.500 services and information types to international ISO and Australian GOSIP standards.

Features of the DX500 OpenDirectory database management system include active query optimisation, on-line automatic backup and multi-site database replication and distribution. It also supports key-word searches across all data in the database, with less than one-second retrieval time, said officials.

Datacraft (03) 727 9111

Print Server

Advanced Components and Peripherals (AC&P) has released the ExtendNet MPX for Token Ring, a high performance multiport Token Ring print server which is able to address up to four protocols simultaneously and transfer print jobs up to 30 times faster than traditional PC-based print servers.

The new product concurrently supports NetWare, LAN Manager, TCP/IP and LAN Server networking protocols.

It allows parallel printers to be located up to 150 metres away from the print server and serial printers to be located up to 60 metres from the print server, while still maintaining full parallel data rates. It can connect up to four printers via high-speed parallel outputs. incorporates Flash memory for easy upgrades, and can be configured and monitored from one location.

AC&P (03) 720 4344

SIMULTANEOUS WAN/LAN ANALYSIS

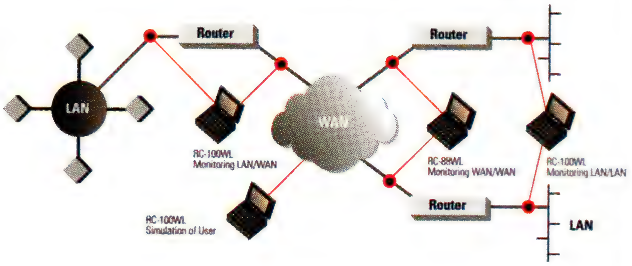

INTERFACES

Token Ring, Ethernet, V.24, RS422, V.35, X.21, E1 (G.703/G.704)

WAN PROTOCOLS

ASYN, SYNC, X.25, SNA/SDLC, Frame Relay, PPP, Encapsulated LAN Protocols

LAN PROTOCOLS

LLC, Mac, NetBeui, NetBios, SNA, IPX, NCP, TCP/IP, SNMP, VINES, Appletalk, DecNet, XNS, OSI, NFS, X-Windows

PORTABLE WAN/LAN ANALYSER,
CALL FOR A DEMO OR INFORMATION KIT TODAY

INS

Information Network Solutions
Unit 3a, 3 Lanceley Place, Artarmon,
NSW 2064. Fax: (02) 906 6327
Phone: Sydney. (02) 906 6335.

CAN YOU GET 2.048MBPS DATA OVER YOUR TELEPHONE NETWORK?

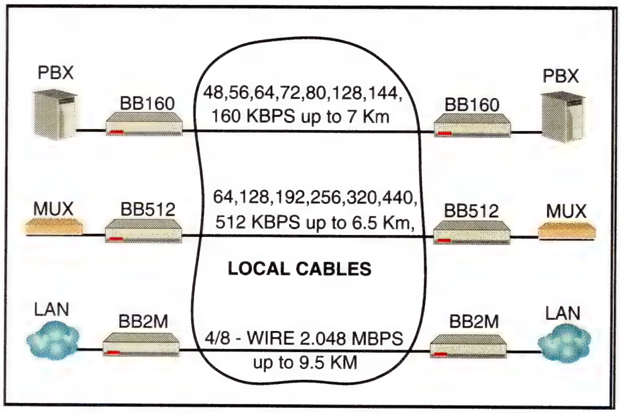

We have a range of high speed, low cost baseband modems capable of speeds up to 2.048MBPS, full duplex point-to-point connection on typical subscriber loop up to 9.5km. High speed baseband modemare is an economical choice for high speed data rates over limited distances.

INS

Information Network Solutions
Unit 3a, 3 Lanceley Place, Artarmon,
NSW 2064. Fax: (02) 906 6327
Phone: Sydney. (02) 906 6335.

INTERNATIONAL DIRECTORIES

Jaeger+Waldmann

New 1994 J+W International Fax Directory

7 volumes, 4 million contacts in 198 countries.
Now available ex-stock Sydney.

International **Telex, Business, Banking** and **Travel** directories (all data also on CD-Rom).

SE Asian directories in English:
Japan, Singapore and Taiwan Yellow Pages.
China Business Phone, Fax and Telex Books.

For all your overseas directories needs, contact sole Australian agents:

OTT TELECOMMUNICATIONS

Melbourne Tel: (03) 795 7426, Fax: (03) 701 1974
Sydney Tel/Fax: (02) 358 6964
Queensland Tel: (074) 47 1263, Fax: (074) 49 8000

MAKING SENSE OF BUSINESS COMMUNICATIONS IS OUR BUSINESS

- SALES, INSTALLATIONS AND SERVICE OF KEY TELEPHONE SYSTEMS
- SALES, INSTALLATIONS AND SERVICE OF P.A.B.X. SYSTEMS
- CATEGORY 5 UTP WIRING NETWORKS FOR BOTH VOICE AND DATA SYSTEMS INCLUDING OPTICAL FIBRE BACK-BONES
- TELEPHONE HANDSETS, FAX MACHINES
- ANSWERING MACHINES, POCKET PAGERS

I.C.C.
Integrated Communications Cabling Pty. Ltd.

Suite 3, 26-28 Aberdeen Road, Macleod, VIC 3085
Tel:(03) 4573600 Fax:(03) 4572782

How can you reach your target audience and backup your sales pitch?

Order customised reprints of your article
or review from
AUSTRALIAN COMMUNICATIONS,
specially designed to
let potential clients read all about
you!

For more information contact:
Head Office : (03)-646 7133 Facsimile : (03)-646 7420

HWM
53 JOHNSTON STREET
PORT MELBOURNE, VIC. 3207

AUSTRALIA'S REPRINT SPECIALISTS

HOWE WILDE MANAGEMENT

Wilpol Pty. Ltd. A.C.N. 006 755 139

CONSULTANTS CORNER

Lane Telecommunications

Specialist Information Technology and Telecommunications consultants, operating throughout Australia and S.E. Asia

- Corporate Communications Policy development and business planning
- Strategic IT&T planning and feasibility studies
- LAN and WAN network design and implementation
- System design, specification and implementation
- Network management strategies and implementation
- Outsourcing strategies and facilities management
- Project management
- Commercialisation of telecommunication service providers
- Benchmarking

COVERING DATA • VOICE • VIDEO • IMAGE • TEXT • RADIO • CABLING SYSTEMS

	Sydney	Adelaide	Canberra	Brisbane
Telephone:	(02)415 1099	(08) 364 5225	(06) 248 9377	(07) 236 4585
Facsimile:	(02) 415 1299	(08) 364 5335	(06) 248 9993	(07) 236 4581

AAS

AMOS AKED SWIFT PTY.LTD.
Telecommunications Consulting Engineers

The largest . . . . .

. . . . . AAS has on-staff expertise to address the largest and most diverse corporate telecommunications assignments

And the smallest . . . . .

. . . . . and also offer tailor-made Packaged Services

AAS Consulting Services . . . . .
solutions to problems — directions for growth

Level 3, 541 Kent Street , Sydney NSW 2000.
Telephone: (02) 264 9640 Facsimile: (02) 264 9868

RAM
COMMUNICATIONS CONSULTANTS
OF AUSTRALIA PTY. LIMITED

A COMMUNICATIONS CONSULTANCY

- ■ Mobile Communications
- ■ Radio Engineering
- ■ Telecommunications
- ■ Data Communications
- ■ Systems Integration
- ■ Communications Facility Planning

KEEPING PACE WITH CHANGING TECHNOLOGY

200 Arden Street, North Melbourne VIC 3051
Telephone: (03) 321 0183
Facsimile: (03) 321 0184

PROTOCOL CONSULTING

"Recruitment consultants to the Telecommunication industry"

Let's talk about the following opportunities:-

- **Software Design and Development**
- **Hardware Engineering**
- **Systems Engineering**
- **Quality Assurance**
- **RF Design and Engineering**
- **Technical Support**

We know you would like more information - so contact *Marisa Lo Certo* **NOW** for a profile of our services and details of the jobs above.

Tel: **(02) 410 9946** Fax: **(02) 410 9947**
Address: Level 6, 10 Help St. Chatswood NSW 2067

Paul Budde
Communication

Some things in life are FREE!

We KNOW the telecommunications industry and we understand your needs! To prove it we'll give you 3 issues of our Telecommunications Management and Marketing Newsletter . . . FREE. Consider this our way of extending our services to you for the next 3 months!

CALL NOW!

Telephone: (049) 988 144
Facsimile: (049) 988 247
2643 George Downes Drive, Bucketty NSW 2250

CONSULTANTS CORNER

is featured in every issue of

AUSTRALIAN
communications

To advertise call
Dawn Switzer on (02) 264 2200

ATUG
AUSTRALIAN TELECOMMUNICATIONS USERS GROUP

Editorial by Wally Rothwell

Austel's 'DMF' is Important to Us All

In early May, the Government passed an amendment to the *Telecommunications Act* which clarifies the operation of the dominant carrier pricing rules and gives Austel the necessary power to disallow anti-competitive tariffs.

Austel, as the telecommunications regulator, will have to oversee these new pricing rules and, in order to do so must develop a Decision Making Framework (DMF). This DMF must be clear to all parties, and it will assist Austel to form an opinion as to whether tariffs or other charging options, offered by the dominant carrier (Telecom) are discriminatory or might have an anti-competitive effect.

ATUG has had a significant input to the legislation amendment and is satisfied with its outcome. We have also been brought into close consultation about the DMF by Austel, and are equally satisfied with the way Austel is going about the job of developing an equitable and sustainable decision-making process.

Methodology

To assist it in this process, Austel has engaged very well qualified consultants in a number of areas, e.g. to establish a definitive practical methodology to deal with the question of whether a carrier 'is in a position to dominate' for provision of a telecommunications service.

That question of dominance is very important to the thrust of our whole pricing regime. If a carrier is dominant, it has very little flexibility in the pricing of its services. A non-dominant carrier can discriminate tariffs to gain competitive advantage. That is, it can offer discounts to certain customers but need not offer the same discount to others for the same service. The 1991 Act intended this to be so, in order to provide a pro-competitive environment, and it has worked well so far.

But dominance is only one component, or module, of the Decision Making Framework. Other modules will address such matters as price discrimination, predatory pricing, cost-justified pricing, damage to competition and competitive market structures.

Important Foundation

The task that confronts Austel is an important one. This is because, in association with the amendments to the Act, it will provide the basis for price competition for the foreseeable future, or at least until mid-1997, when full blown competition should be upon us.

Austel is going to have to determine first what the 'markets' are before it can begin to decide whether a carrier is dominant in a particular market. It will need to define what is 'commercially sustainable competition' and what is the legal definition of discriminatory pricing. There needs to be a distinction between equality of opportunity and of outcome. And while cost-based tariffing is something that ATUG has pursued assiduously, how close can a carrier get to this ideal and how can that be tested?

In short, this is not a simple task. Along with its use of expert consultants, Austel has called for industry submissions. As this issue goes to press, ATUG will be making its initial submission, due on May 29.

The consultants will digest the submissions they receive and report to Austel in mid-June, with a 'first-cut' DMF being provided at the end of June. Subsequently, Austel will conduct an open seminar in Melbourne on 15 July and hopes to have the DMF in operation during August.

Interim Arrangements

In the interim, there is a flood of new Telecom tariffs which we expect to be filed with Austel this month.

Austel is developing arrangements for the interim period between now and August whereby the dominant carrier may file tariffs which will be subject to a 'sanity test' which is in line with what Austel believes to be the intent of the legislative amendments.

ATUG expects that the revised in-dial tariff, which it has been negotiating with Telecom for nearly two years, will be one of those released — possibly as a Flexi-Plan — under these interim arrangements.

ATUG has continued to press for reasonably predictable tariffing arrangements so that businesses can plan with some degree of certainty. So the DMF is perhaps the most important tariffing development in our evolving competitive environment to occur this year.

This new Decision Making Framework will provide Austel with a reasonable certainty about whether or not a tariff is anti-competitive. The DMF is vital to the planning process of our members and we will stay across it in close detail.

Wally Rothwell

Wally Rothwell
Executive Director

AUSTRALIAN TELECOMMUNICATIONS USERS GROUP

ATUG Launches Something Extra

Members may be able to reduce their telecommunications costs by up to 30% following the introduction of a new savings scheme by ATUG.

The scheme, launched at ATUG'94, is called the 'ATUG Extra Corporate Discount Plan.' By taking advantage of the liberalised telecommunications marketplace under this scheme, ATUG is able to provide discounts on telecommunications services and other corporate services.

According to ATUG's Deputy Executive Director, Richard Allen, this means that users will be able to achieve discounts greater than they could obtain in their own right. "While the discounts will vary from organisation to organisation, discounts against certain Telecom tariffs can be as high as 30%," he said. Examples of the savings which users may be able to achieve via this scheme are:

- For a one-person professional office, with combined phone, fax and mobile costs of about $200 each month, the savings under the ATUG Extra plan would be in excess of $600 per annum;
- For a 13-person industry association, with combined phone, fax and mobile costs of $1,300 per month, the savings would be approximately $3,200 per annum;
- For a small business with 20 personnel, and monthly telecommunications costs of around $2,800, the savings would be around $7,500 per annum;
- For a 200-staff subsidiary of an overseas multinational, with a monthly bill of around $14,000, annual savings can be achieved in the order of $27,000.

As well as reduced telecommunications costs, the scheme also provides discounts on office equipment, travel, wine and insurance. The savings under the ATUG Extra plan are in addition to existing ATUG membership benefits like discount rates to the annual ATUG convention and ATUG seminars.

"Historically, ATUG's activities have achieved much for the marketplace at large. The ATUG Extra plan is viewed as one way to reward member organisations for their loyal support," Allen said. Access to the Corporate Discount Plan is available to all paid-up members.

ATUG's Plan For The Future

ATUG's mission is to ensure that Australian users get the widest range of best quality communications services at the lowest prices via healthy competition.

With this in mind, ATUG Board members and ATUG personnel get together for an annual Strategic Planning Retreat weekend to discuss areas of concern to members.

The broad general objective is to ensure that ATUG does not become a victim of circumstances beyond its control, but, rather, by establishing its long-term strategic vision, takes advantage of each opportunity to achieve desired outcomes. ATUG has outlined a number of issues to ensure that users continue to get the best deal in the future, including the following policy issues:

- Ensuring that open competition is introduced in 1997;
- Ensuring that carriers adopt appropriate tariff arrangements;
- Revising the current 'Basic Carriage Services' and 'Higher Level Services' definition in the *1991 Telecommunications Act* to reflect the desired outcomes by 1997;
- Ensuring that an Australian Personal Communications System fully complies with a single international data protocol;
- Enhancing competition by introducing portability, allowing users to change carriers for services while retaining numbers for those services;
- Encouraging competition by using foreign satellites for the provision of domestic satellite services;
- Ensuring that rural and remote users are provided with access to telecommunications services as good as those available to urban users;
- Directing the outcome of the 'Reach Out' trials (this is the Extended Untimed Local Call Area project) to the best long term advantage of users; and
- Ensuring the availability of the widest range of high quality international services at the lowest prices via healthy competition.

ATUG's policies are published in the 1993 Annual Report. Members are encouraged to read and comment on these to ensure they reflect members' views.

The next Strategic Planning Retreat will take place in October 1994. Members' comments are welcomed before the end of September this year.

Any comments should be faxed to: ATUG 1994 POLICIES on (02) 925 0880.

SYDNEY
Level 11, 80 Alfred Street
Milsons Point, NSW, 2061
Australia
Tel: (02) 957 1333, (008) 22 6281
Fax: (02) 925 0880
MELBOURNE
Unit 3, 8-12 Sandilands Street
South Melbourne, VIC, 3205
Australia
Tel: (03) 690 6395
Fax: (03) 699 9094

ATUG'94 Award Winners

Each year, ATUG takes the opportunity to honour individuals for their contributions to Australian telecommunications. The main award presented is the Charles Todd medal for the 'Communicator of the Year,' named after Sir Charles Todd (one of Australia's great pioneers, who conceived, planned and personally supervised construction of the Overland Telegraph Line, which was completed in 1872).

Michael Tkacz was awarded this year's Communications Management Excellence prize

This year's winners were announced during ATUG'94 in Melbourne last month.

The 1994 'Communicator of the Year' award, sponsored to the value of $5,000 by *The Australian* newspaper, was won by JNA Telecommunications' Chairman, Mr John Almgren, for his outstanding service to the communications industry and the Australian community for over thirty years.

He established the company with his wife in 1960 from a small bespoke communications engineering business and has developed it into a significant Australian public company with an expanding international business.

Another award winner this year was Mr Stuart Corner, publisher of the weekly newsletter *Exchange* and the fortnightly newsletter *Telenews Asia.* Mr Corner was presented with the 'Telecommunications Journalist of the Year' award for 1994, in recognition of the important tasks which journalists play in keeping the Australian public informed about developments within the telecommunications industry.

John Almgren shortly after winning the award for Communicator of the Year

Finally, this year's winner of the 'Communications Management Excellence' award was Mr Michael Tkacz from Curtin University. The runner-up was Mr Frank Vardanega from the Brisbane City Council.

These important awards are sponsored by leading telecommunications consultancy Amos Aked Swift to the value of $3,000 and $1,000 respectively.

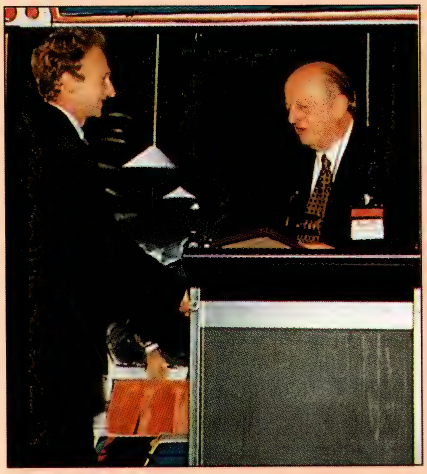

Stuart Corner receiving his award from ATUG Chairman George Maltby

ATUG Salutes AIN Participants

ATUG presented awards to participants in the ATUG Interoperability Network (AIN) to recognise their commitment to the project, which formed part of the ATUG'94 Exhibition.

The AIN was developed to demonstrate simple interconnectivity via FDDI, and provided this year's exhibitors with the opportunity to communicate with other participating ATUG'94 exhibitors and their branch offices, and also to establish international connections via the Internet.

The ATUG Interop. Network Partnership Awards were made to individuals who had shown outstanding commitment and co-operation in making the AIN a success. Winners were Amos Aked Swift's Brian White, Norman Gale of Nortel, Geoff Letts, Bill Hayne and Peter Stevenson from JNA Network Services, Mike Snowdon of the Commonwealth Bank, and Justin Cooke from ATUG.

Companies who participated in the project and demonstrated genuine interoperability received the ATUG Interop. Network Gold Award. They were Novell, Hewlett-Packard, 3Com, SynOptics, Proteon, Cabletron, Nortel, Digital Equipment Corporation and Wellfleet.

ATUG AUSTRALIAN TELECOMMUNICATIONS USERS GROUP

The Value of ATUG Membership

Can you afford to be without the up-to-date knowledge, provided by ATUG membership, on the evolution of convergence in the media, communications and computer industries?

As a member of ATUG you receive:

⇒ One year's free subscription to the monthly ATUG newsletter, 'Newsbrief';
⇒ One year's free subscription to *Australian Communications*;
⇒ Discount rates to the telecommunications event of the year, ATUG's annual four-day conference and exhibition, ATUG '94;
⇒ Discount rates to ATUG seminars and workshops;
⇒ Free admittance to ATUG special member meetings;
⇒ Free assistance on matters concerning telecommunications;
⇒ Representations of your views to Austel and the Government;
⇒ A forum to exchange ideas with your peers in the industry;
⇒ A telecommunications library service;
⇒ Member rates for ATUG products and services;
⇒ Expert legal advice from ATUG's Legal Line service; and
⇒ An overall saving on your next communications bill due to ATUG's lobbying activities.

If you have any queries about your ATUG membership, call ATUG's business manager, Owen Richards, on (03) 690 6395. Members in Queensland may contact ATUG's business manager, Jane Prentice, on (07) 870 8831. Members in Western Australia may contact ATUG's WA business manager, Geoff Groves, on (09) 450 3485.

ATUG Network Management Survey

Copies of the report of ATUG's Network Management Survey are still available — but not for too long! The complete report, including quantitative data, runs to over 40 pages together with well over 60 pages of supporting information and suppliers' responses. The report is available at $125 (including postage and handling) to ATUG members; and at $200 to non-members. For copies of either the complete report or the Executive Summary, please complete the form below:

Please forward ______ (qty) copies of the ATUG Network Management Survey at $125 for each copy ($200 non-members) for which I enclose payment of: $______

Please tick payment method (DO NOT send cash):

☐ **Cheque** ☐ **Mastercard** ☐ **Bankcard** ☐ **Visa**

Cheques should be made payable to ATUG.

Credit card number:

☐☐☐☐ ☐☐☐☐ ☐☐☐☐ ☐☐☐☐

Expires: ________ Cardholder's Signature: ____________________

Cardholder's Name: ____________________

Company: ____________________

Address: ____________________

____

____

OR:

☐ **Please provide a copy of the six page Executive Summary, free of charge to ATUG members only.**

Name: . . . .
Tel: . . . .
Fax: . . . .

Forward to:
ATUG
Network Management Survey
PO Box 357
Milsons Point, NSW 2061

Frame Relay Forum Announces On-line User Services

The Frame Relay Forum has announced two new on-line information services which interested parties can use to obtain and exchange information on frame relay technology and its applications.

These services are free and easy to use. They include an electronic bulletin board called 'The Frame Relay InfoExchange' and a fax-on-demand service known as 'The Frame Relay Facts-by-Fax.'

Network managers have the ability to read and download articles about other organisations' experiences with the technology as well as the ability to interact directly with other users.

On-line Interaction and Documents

Inspired by the growing popularity of on-line services and the use of the Internet to access current information, the Frame Relay InfoExchange has an information library and an interactive message section. The Library contains electronic copies of documents such as:

- An overview of frame relay;
- A bibliography of recent articles regarding frame relay, with complete texts of selected articles;
- Copies of frame relay Forum newsletters;
- A directory of standards related to frame relay, including proposals in the draft stage and full copies of selected standards;
- A directory of Frame Relay Forum member companies, with contact information; and
- Information on upcoming Frame Relay Forum activities.

The interactive message section provides direct interaction between current and potential Frame Relay users. The InfoExchange allows people to ask questions, receive answers, post useful information and browse through messages posted by users and others interested in the technology. For example, an individual considering frame relay might want to compare how one particular application behaves using frame relay instead of multipoint leased lines. By posting a message a user can determine if any other InfoExchange user has already implemented frame relay with that application.

Facts-by-Fax Information

Facts-by-Fax is an expansion of a fax service initiated by the Frame Relay Forum in 1992. The documents available from Facts-by-Fax include a subset of the materials in the library portion of the InfoExchange.

To access this service, users dial Facts-by-Fax from the telephone handset of a standard fax machine and follow the instructions. By responding to the voice prompts with touch-tone entries, they can receive a listing of available documents, or can request one or more specific documents. Facts-by-Fax will then automatically fax the requested information to the user's fax machine.

How to Get On-line

The Frame Relay InfoExchange is operated in cooperation with CompuServe, which is providing access to the InfoExchange free of connection or connect-time charges.

Individuals who have a CompuServe account can access the InfoExchange by typing 'GO FRAME.' They can either enter this exchange as a paying visitor, or they can input their name to register and join the InfoExchange group of free usage.

Those who do not have a CompuServe account and Internet users who wish to access the InfoExchange can obtain information on how to connect by contacting the Frame Relay Forum on (02) 975 2582.

ATUG Legal Line Service

Members should take advantage of ATUG's Legal Line Service. The service is run by ATUG director and chairman of ATUG's Legal sub-committee, Gerald Wakefield.

Members with legal queries regarding international telecommunications law and regulatory policy; radiocommunications; broadcasting and satellite law and policy; computer technology law; and intellectual property law are advised to seek this communications legal advice service by calling ATUG on (02) 957 1333.

Telecommunications Act on Disk

The new EIS Electronic Law Book now includes all the updates to the Telecommunications legislation along with Broadcasting and Radiocommunications legislation.

The publication, on a floppy disk, is available to ATUG members at the special price of $150 including postage and packaging — a saving of $50. A subscription including a further two updates is also available for the special price of $295 — a saving of $100.

To ensure that you have the legislation at your fingertips at all times, take advantage of this special offer.

Contact ATUG on (02) 957 1333 for details of how to order.

June

14-16 **Marketing, Implementing and Distributing GSM**, Regents Park Marriott, London. This event features speakers from the world's leading GSM equipment vendors and network operators. Delegates will hear both how to overcome technical challenges posed by implementing a network in a fast-changing technological environment, and how to market a GSM service while still capitalising on existing markets, and how to develop value added services to generate new revenue and compete with rival companies' strategies. Fee: Conference only £934; Optional one-day seminar, 'Rolling Out GSM,' £581; Conference and seminar £1,339. Enquiries — IIR Conferences Tel: +44 71 412 0141 Fax: +44 71 412 0145.

14-15 **Demystifying LANs**, Southern Cross Hotel, Melbourne. This seminar aims to give delegates a thorough overview of LAN technology and a solid foundation in networking fundamentals. It covers topologies; cabling and access methods; protocols and architectures; the role of hubs, bridges, routers, gateways, network operating systems and the use of network management tools. The issues of security and reliability will also be discussed. Fee: $994. Enquiries — MTE Tel: (02) 261 5555 Fax: (02) 261 5959.

22-24 **Wireless Local Loop**, Kowloon Shangri-La, Hong Kong. The Asian region represents a huge potential market for the implementation of wireless local loop technology, as it provides the ability to provide sophisticated telephone services to countries with underdeveloped telecommunications infrastructures without the need for enormous capital investment. This conference looks at a broad range of issues involved in implementing wireless local loop systems, including regulation, planning and finance, the role of the technology in growing markets, and comparing different technologies such as DECT, CT2 and VSAT services. Fee: $US1,895. Enquiries — IIR Conferences Tel: +852 549 5618 Fax: +852 547 3836.

29-1/7 **PTC Mid-Year Seminar**, Regent Hotel, Sydney. The theme of the mid-year conference of the Pacific Telecommunications Council will be 'Bridging National and Business Cultures: Strategic Alliances in Communications.' Enquiries — PTC Sydney Conference Tel: (03) 370 6470 Fax: (03) 370 6452.

July

1-5 **Telecom China '94 Telecommunications Equipment Exposition**, Guangzhou International Exhibition Centre, Guangdong Province. This exhibition provides an opportunity for overseas companies to gain wide exposure in the lucrative and fast growing Chinese market. Enquiries — Union Fair and Trade Development Tel: + 86 20 888 2823 Fax +86 20 884 9074.

3-6 **International Telecommunications Society 10th Annual Conference**, Sydney Convention Centre, Darling Harbour, Sydney. This year's conference will look at the convergence of telecommunications, mass media and computer technologies brought about by developments in intelligent interconnected systems. The event will focus specifically on the areas of broadband infrastructure and user requirements; regulation of the industry; competition issues; and the contribution of telecommunications to economic growth. Enquiries — Donald Lamberton, ITS Tel: (06) 249 3884 Fax: (06) 249 0312.

25-26 **Intelligent Networks**, Hyatt Kingsgate, Sydney. The growth in Intelligent Network services in Australia is just beginning to take off, with suppliers of switching, computer equipment and software predicting a boom in demand for new services. This event features speakers from the carriers, equipment vendors, leading consultancies, regulators and overseas operators, and will explore opportunities using IN solutions, evolving network architecture standards, and will also feature overseas case studies. Fee: $1,395. Enquiries — IIR Conferences Tel: (02) 954 5844 Fax: (02) 959 4684.

25-27 **Remote LAN Access**, Sydney Hilton. This conference will give delegates an overview of new remote access technologies to help them better define networking strategies and optimise their network. Emerging trends such as client/server computing will be examined, and the programme also looks at how to maintain a secure environment while still permitting essential remote links. A separately bookable one-day workshop focuses specifically on implementation issues. Fee: Two-day conference $1,395; workshop only $895; conference and workshop $1,995. Enquiries — AIC Conferences Tel: (02) 210 5777 Fax: (02) 221 7773.

27-28 **Electromagnetic Compliance and Interference**, Golden Gate Park Plaza, Sydney. In 1996 the European Union will introduce a new and comprehensive Electromagnetic Compliance standard which will affect all manufacturers of goods and components which are susceptible to electromagnetic emittance or interference. To remain internationally competitive, local manufacturers must comply with these new standards, and to facilitate this, Australia is currently developing standards for EMC/EMI compliance. This conference will update delegates on the latest state-of-play, including how EU standards are being implemented, the impact on exports, the Spectrum Management Agency's testing requirements, and litigation issues. Fee: $1,095. Enquiries — IBC Conferences, Tel: (02) 319 3755 Fax: (02) 699 3901.

27-29 **Multimedia '94**, Hotel Nikko, Sydney. This event will give a broad overview of what is meant by the term multimedia, and what can be expected from the imminent boom in new applications which take advantage of the ability to incorporate sound, video, graphics and text. There will be leading industry speakers, and the conference will also look specifically at multimedia for business applications. Fee: Three days $1,895; two days $1,395. Enquiries — IIR Conferences (02) 954 5844 Fax: (02) 959 4684.

August

2-4 **Comtec '94**, Adelaide Exhibition Hall and Convention Centre, Adelaide. This event is expected to attract over 15,000 people, and will feature exhibitors such as Novell, Microsoft, Apple and Intel. There will be free seminars on specific products, while the conference will address the broader issues of today's computing industry. Enquiries — (08) 364 4723.

3-5 **Utility Competition and Regulation Summit '94**, Savoy Park Plaza, Melbourne. This conference features top level speakers from government, industry and regulating organisations, and will concentrate on the implications of Australia's emphasis on the development of a competitive environment and opening up monopolistic utility markets. There will also be a special workshop entitled 'Benchmarking Utility Performance' on the third day. Fee: Two-day conference $1,495; workshop only $895; conference and workshop $1,995. Enquiries — Karen Monks Tel: (02) 210 5777.

3-6 **TUANZ '94**, Ellerslie Function Centre, Auckland. TUANZ'94, the annual conference and exhibition of the Telecommunications Users Association of New Zealand will have the theme 'Convergence in Action,' and will focus this year on the convergence of computing and telecommunications. For the first time the conference will include separate streams focusing on fields of interest, and there will be a seminar on Technology in Education. TUANZ organisers say they expect the exhibition space for the 1994 event will be expanded. Enquiries — TUANZ Tel: +64 9 488 1602 Fax: +64 9 489 9515.

September

25-28 **LETA 94**, Adelaide Convention Centre, Adelaide. This conference, Learning Environment Technology Australia, will explore, explain and demonstrate the effects of new and forthcoming technologies on learning and the way it is undertaken. As well as teaching issues there will be a focus on building design and adaptation to meet new requirements. The event features a wide range of speakers from both Australia and overseas. Enquiries — Tel: (08) 226 1266 Fax: (08) 226 1583.

26-27 **Communications Research Forum 1994**, Hyatt Kingsgate, Sydney. This event, targeted at the research community, looks at current and recent developments relevant to emerging public policy issues in the communications arena. The focus of the two-day conference will be on telecommunications, broadcasting and the radio spectrum, looking specifically at policy-relevant work into economic, sociological, cultural and legal implications. Enquiries — Vickie Richardson Tel: (06) 274 6846 Fax: (06) 274 6816.

October

25-28 **Communications India**, Pragati Maidan, New Delhi. This trade show will display a full range of telecommunications hardware and software in the areas of voice, data, text, visual display and radio. An associated technical conference will examine cellular mobile, networking, rural telephony and network privatisation. There will also be keynote presentations by the ITU, UNIDO, the World Bank, the Indian Department of Telecommunications and the Telecom Equipment Manufacturers Association of India. India is one of the fastest growing telecommunications markets worldwide. Enquiries — Exhibitions India Tel: +91 11 462 2710 Fax: +91 11 463 3506.

25-27 **Microwaves '94**, Wembly Conference & Exhibition Centre, London. This conference promises to bring together the largest gathering in Europe of microwave experts engaged in research and development for both telecommunications and civil and military systems. The program will cover the entire microwave and millimetre wave spectrum, with supporting RF and software technologies. Enquiries — Nexus Business Communications Tel: +44 322 660 070 Fax: +44 322 667 633.

December

5-7 **ATNAC '94**, Hilton Hotel, Melbourne. This event incorporates the Multimedia Communications, Applications and Technology Workshop (MCAT), the Australia Broadband Switching & Services Symposium (ABSSS), and the Australian Teletraffic Research Seminar (ATRS). The theme of the conference will be 'Meeting User Needs,' and keynote speakers will include Phillip Dodds, President of the Interactive Multimedia Association of the USA and Leonardo Chiariglione, convenor of the Motion Picture Experts Group (MPEG). The registration fee for the three-day conference is $450. Enquiries — Margaret Keegel Tel: (03) 903 2808 Fax: (03) 903 2805.

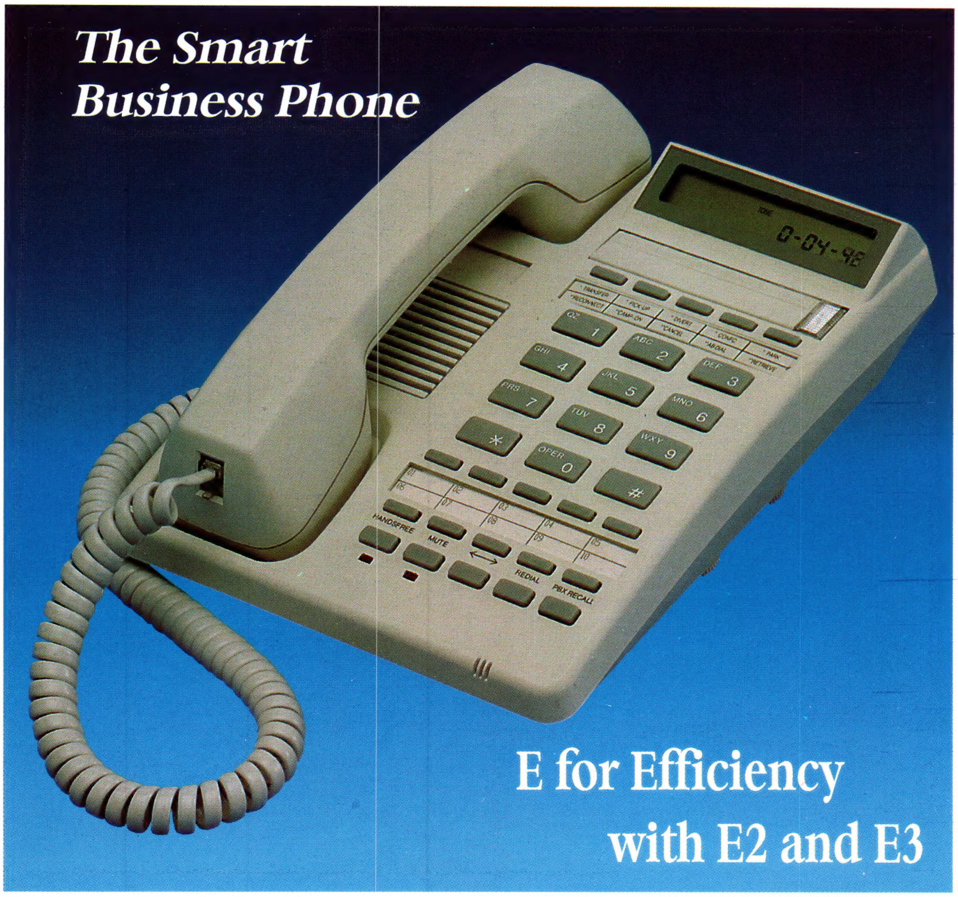

IQTEL E3

- Liquid crystal display
- Clock
- Handsfree / conference operation
- Adjustable earpiece and handsfree speaker volume
- Line powered
- Memory capacity 20 Telephone numbers or PABX commands including 10 numbers with non-erasable Eeprom insertion
- Soft touch number key pad
- Colour Grey

IQTEL E2

On Hook / Monitor Dialling Version

The following feature is optional

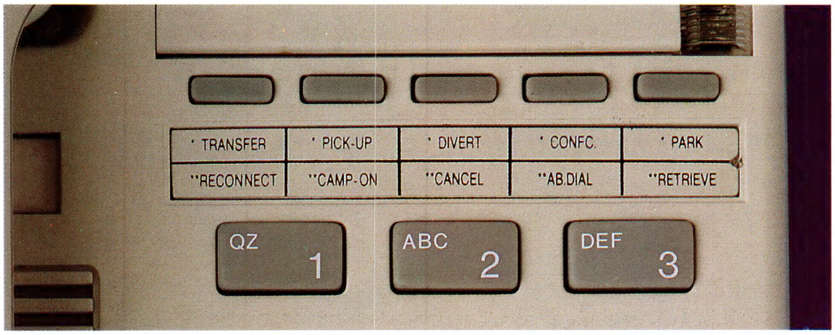

These phones not only look good, they will help you through the business day.

Easy to identify, easy to use PABX or carrier network facilities.

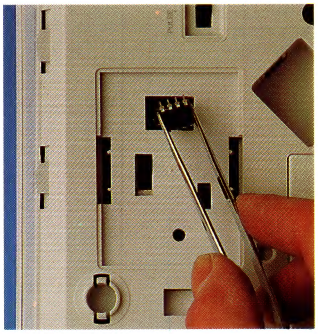

We fit this optional, specially masked Eeprom, permanently pre- programmed with your PABX or carrier network facilities

Maybe it is just checking the time or a telephone number being dialled from the display, talking handsfree or dialling on hook, phoning a colleague with one key touch or instructing the PABX or carrier network facilities to "call forward" your calls by one touch of a memory key. Whatever your job is, you will find E2 or E3 a very dependable workmate.

E2 and E3 can be purchased with customised pre-programming of PABX or network services such as Centel/Customnet/Centrex facilities as illustrated on the left of this page. The pre-programming is on a permanent basis, and will result in all staff being able to easily identify and use the PABX or Carrier facilities. The training time of staff in telephone feature use can be greatly reduced or even eliminated with E2 and E3.

E2 and E3 are not only business phones, they are ideal for home use when purchased with Telecom's "Easy Call" facilities pre-programmed into Eeprom memory for one touch activation.

PLEASE NOTE: E2 & E3 MODELS AVAILABLE DECEMBER 1993.

IQ INTERQUARTZ® – The Manufacturers and Distributors

INTERQUARTZ (A'ASIA) PTY. LTD

Melbourne Phone: (03) 457 5000
Fax: (03) 458 4049

Sydney Phone: (02) 875 4011
Fax: (02) 875 4146

Brisbane Phone: (07) 808 8008
Fax: (07) 808 8080

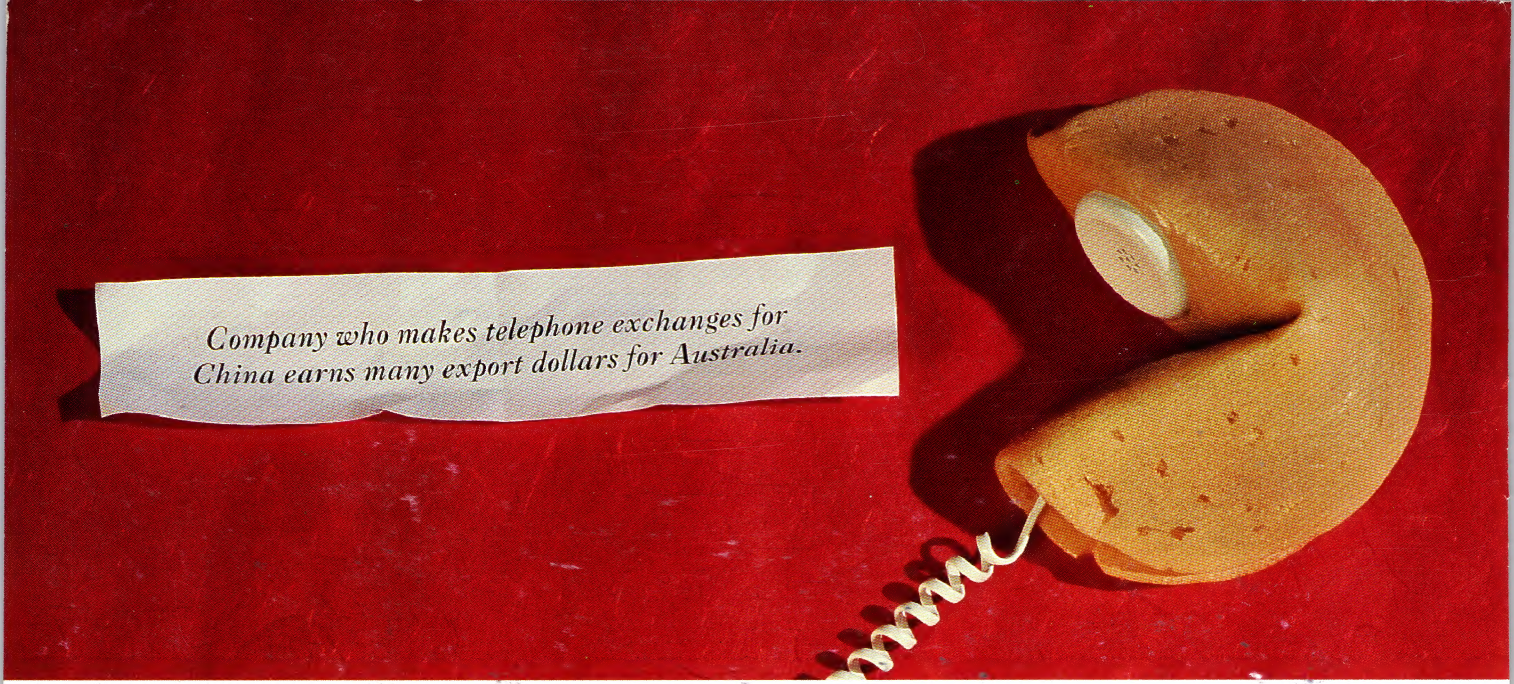
Company who makes telephone exchanges for China earns many export dollars for Australia.

Alcatel Australia has cracked the fiercely competitive market of China. We've had the good fortune of winning contracts for our System 12 digital telephone exchanges in several regions. Sales to China of this switching system are earning Australia millions of dollars in exports. Add this to the $100 million invested in Australia to design and manufacture these systems and we're talking big bikkies. At the same time we're helping the 1.1 billion Chinese get better connected. And we're certainly no strangers to the global marketplace. Total exports for last year amounted to $200 million. This has helped us become Australia's leading exporter in the telecommunications industry. And so long as that's the way the cookie crumbles, there'll be more export earnings in the future. Alcatel Australia. We're Australia's largest communications company because we know how to compete worldwide.

ALCATEL
AUSTRALIA

LBC&M ALC 0106